AD_____

Award Number: W81XWH-06-1-0003

TITLE: Intra-Operative Dosimetry in Prostate Brachytherapy

PRINCIPAL INVESTIGATOR: Ameet Jain

CONTRACTING ORGANIZATION: Johns Hopkins University
                                         Baltimore, MD 21218-2686

REPORT DATE: November 2006

TYPE OF REPORT: Annual Summary

PREPARED FOR: U.S. Army Medical Research and Materiel Command
                       Fort Detrick, Maryland 21702-5012

DISTRIBUTION STATEMENT: Approved for Public Release;
                                       Distribution Unlimited

The views, opinions and/or findings contained in this report are those of the author(s) and
should not be construed as an official Department of the Army position, policy or decision
unless so designated by other documentation.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
01-11-2006	Annual Summary	15 Oct 2005 – 14 Oct 2006

4. TITLE AND SUBTITLE

Intra-Operative Dosimetry in Prostate Brachytherapy

5a. CONTRACT NUMBER

5b. GRANT NUMBER
W81XWH-06-1-0003

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Ameet Jain

E-Mail: jain@cs.jhu.edu

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Johns Hopkins University
Baltimore, MD 21218-2686

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES
 Original contains colored plates: ALL DTIC reproductions will be in black and white.

14. ABSTRACT

Favorable outcome in prostate brachytherapy critically depends on the accurate placement of radioactive sources in their planned locations. Unfortunately, there is variety of mechanical factors that cause the seeds to divert from their planned locations. While this problem has been known to brachytherapists, current technology does not allow for reliable localization of the implanted sources, thereby prohibiting the prediction and modification of seed distribution intra-operatively. The Research Objective of the proposal is to develop and evaluate ex-vivo a method for intra-operative localization of the implanted seeds in relation to the prostate, to facilitate in-situ dosimetric optimization and exit dosimetry. In particular, we will: [1] Registration of Ultrasound to Fluoroscopy (RUF): Develop methods for reconstruction of seed implants from X-ray fluoroscopy and spatially register them to the prostate anatomy identified in TRUS [2] System Integration: Integrate the above methods in a software package and link it with the FDA-approved CMS Interplant® prostate brachytherapy system to enable in-situ dosimetry calculation [3] Experimental Validation: Evaluate the performance of the RUF system on phantoms and pre-recorded patient data.

15. SUBJECT TERMS
 Prostate Brachytherapy, X-ray reconstruction, C-arm, TRUS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON USAMRMC
a. REPORT U	**b. ABSTRACT** U	**c. THIS PAGE** U	UU	92	**19b. TELEPHONE NUMBER** *(include area code)*

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Table of Contents

# Progress Report Summary
Project Year 1 (2005 –2006)

A   INTRODUCTION

For several decades, the definitive treatment for low and medium risk prostate cancer was radical prostatectomy or external beam radiation therapy, but low dose rate permanent seed brachytherapy (shortly brachytherapy) today can achieve equivalent outcomes. Brachytherapy, if accurately executed, can achieve a sharp demarcation between the treated volume and healthy structures, and thereby achieve superior tumor control with reduced morbidity. In contemporary practice, however, faulty needle and source placement often cause insufficient dose to the cancer and/or inadvertent radiation of the rectum, urethra, and bladder. Another fallacy of the current implant techniques is that reliable and accurate exit dosimetry is not possible. The contribution of the proposed research will be making C-arm fluoroscopy available for safe, simple, and robust intra-operative localization of brachytherapy sources relative to the prostate. We will develop a method for the registration of ultrasound to fluoroscopy (RUF), to fuse TRUS (Trans-rectal ultrasound can view the prostate but not the seeds) with C-arm fluoroscopy (which is capable of viewing the seeds but not the prostate). This feature will allow for dosimetric optimization of the prostate brachytherapy implants and exit dosimetry before the patient is released from the operating room; thereby enabling significant improvement on current clinical practice. A further promise is that fluoroscopy-based exit dosimetry may obviate CT-based post-implant dosimetry.

B BODY
B.1   Brief System Concept

The system concept for registration of ultrasound to fluoroscopy (RUF) is summarized in Figure 1. The fluoroscope is calibrated and corrected from image distortion pre-operatively. The implant procedure starts as usual: TRUS is used to guide each individual needle and a C-arm placed over the patient's abdomen. The C-arm is tracked with an X-ray fiducial system called FTRAC that is composed of optimally selected polynomial space curves. The fiducials are mounted rigidly to the TRUS frame in the
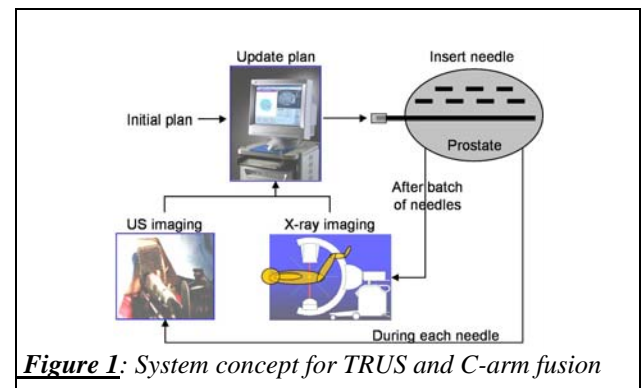
Figure 1: *System concept for TRUS and C-arm fusion*

field of view in a known calibrated pose relative to the TRUS, thereby providing spatial registration between the C-arm and TRUS. Upon implanting a batch of needles (typically a row of needles), we collect a set of TRUS and C-arm images. The locations of the implanted seeds are recovered from the C-arm fluoroscopy images with the use of a network flow based method called MARSHAL that provides seed segmentation, matching, and reconstruction method. Then the cloud of seeds is superimposed over the spatially co-registered TRUS images. The 3D dose distribution is rapidly calculated from the union of the already and yet to be implanted seeds. The dose distribution is analyzed with tools currently available in the brachytherapy system used. Then the implant plan can be optimized to account for discrepancies from the ideal dose distribution. The procedure continues with the next batch of needles in the cycle described above. After the last needle, a complete dosimetry check is performed, which provides a final opportunity to patch up any cold spots with additional seeds.

B.2   Proposed Statement of Work

We proposed to develop a method for the registration of ultrasound to fluoroscopy (RUF) to allow for intra-operative dosimetry in prostate brachytherapy and prototype mathematical algorithms (Aim-1), integrate them with an existing FDA approved prostate brachytherapy system that provides dosimetry analysis (Aim-2), and evaluate the system experimentally on phantoms and pre-recorded patient data (Aim-3). Algorithmic design (Aim-1) and experimental evaluation (Aim-3), will progress hand in hand. System integration (Aim-2) will be
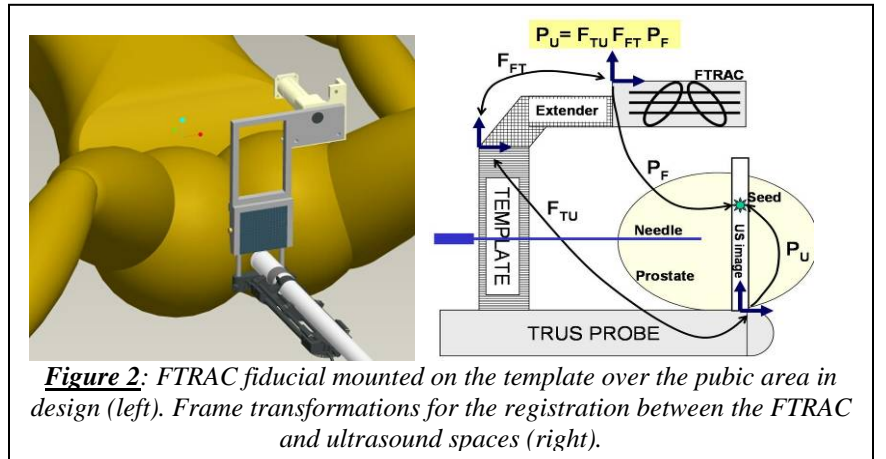
performed immediately when a workable subset of RUF package becomes available from Aim-1 and again revisited towards at the end of the project. Therefore, the timeline will be somewhat non-linear. The detailed statement of work was as follows:

Aim-1: Registration of Ultrasound to Fluoroscopy (RUF): Develop a methods for reconstruction of seed implants from X-ray fluoroscopy and spatially registering them to the prostate anatomy identified in TRUS

Aim-2: System Integration: Integrate the above methods in a software package and link it with the FDA-approved CMS Interplant® prostate brachytherapy system to enable dosimetry calculation

Aim-3: Experimental Validation: Evaluate the performance of the RUF system on phantoms and pre-recorded patient data. (Neither of which require an IRB approval)
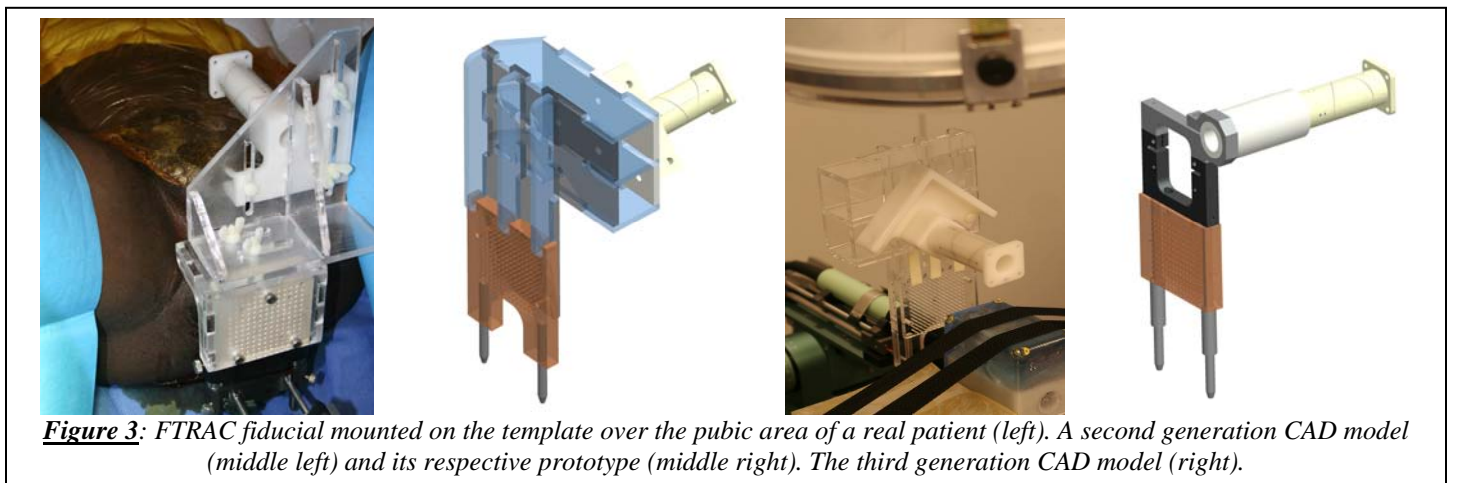
B.3 Progress Report for First Year

<u>In Aim1</u>**:** the three main aspects were (a) Registration of X-ray to ultrasound; (b) C-arm tracking; and (c) Seed reconstruction.

(a) The FTRAC fiducial [12] was rigidly attached to the implant template with a precision-machined mechanical connector in a known position, while the template is already registered to TRUS. The spatial relationship between the FTRAC and TRUS is explained in Figure 2. Let P_F be a point reconstructed in FTRAC space, P_U be the same point in ultrasound space, F_{FT} be the transformation between FTRAC and template, and F_{TU} be the transformation between template and

Figure 2: FTRAC fiducial mounted on the template over the pubic area in design (left). Frame transformations for the registration between the FTRAC and ultrasound spaces (right).

ultrasound. Then the relationship between P_U and P_F can be expressed as $P_U = F_{TU} F_{FT} P_F$. A prototype was manufactured and tested on real patients for dimension verification and workspace feasibility, as shown in Figure 3. Using this data, a new design was created and fabricated, as shown in the middle two images of Figure 3. A calibrated Optotrak (an infrared optical tracker) pointer was used to digitize various lines, which were then used to register the FTRAC fiducial to the insertion template. The accuracy of this registration process was 0.22

Figure 3: FTRAC fiducial mounted on the template over the pubic area of a real patient (left). A second generation CAD model (middle left) and its respective prototype (middle right). The third generation CAD model (right).

mm. Moreover, the attachment was tested for rigidity with respect to sterilization. The pre and post-sterilization registration produced the same values (within the accuracies of measurement), indicating that the fixture is robust to the sterilization process. However, two problems were noticed: (i) drilling of the needle insertion holes had an error of 1º, leading to a 2 mm offset for the seed deposition; and (ii) the weight of the design was not

fully desirable. Hence a new attachment is being designed, which is much lighter and more accurate, as illustrated in Figure 3. The main improvements in the design of the new attachment are: (i) the capability to operate at multiple heights, providing the capability to accommodate obese patients; (ii) superior placement with respect to the patient, providing better workspace optimization; (iii) lighter weight; and (iv) more accurate holes. The final attachment is expected to be ready in the coming few weeks, when appropriate validation experiments would be performed.
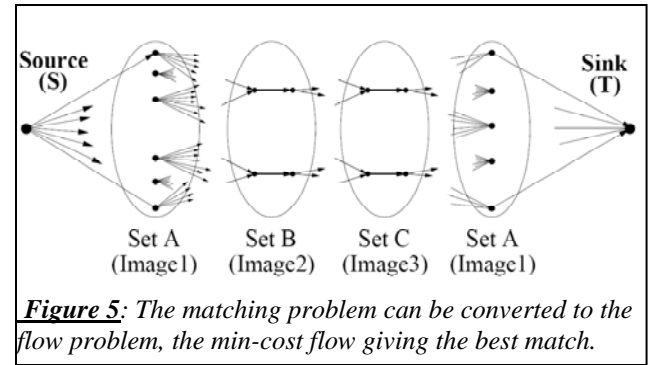
(b) The primary objective here was the improvement in the robustness with which the C-arm can be tracked. The earlier hypothesis was that the accuracy bottle neck for C-arm tracking would be the ability to completely segment the fiducial from the clinical images. We had proposed a
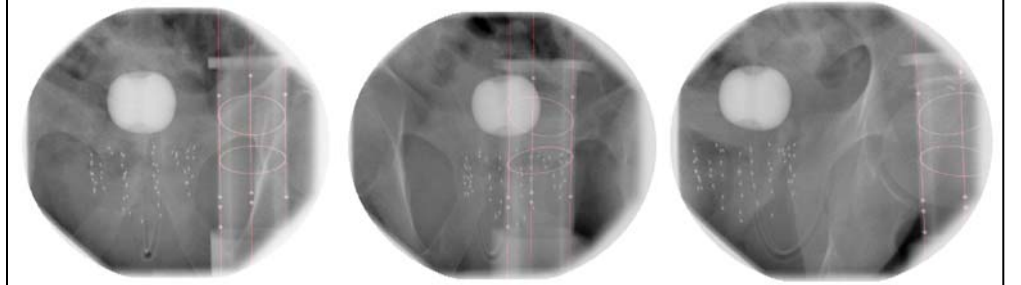
Figure 4: X-ray images with the FTRAC fiducial mounted over the pubic area. It can be noted that the segmentation algorithm is robust to significant variations in the images.

bootstrapping algorithm to tackle this problem, where the segmentation and pose estimation processes were alternated. However, we developed a segmentation algorithm that was sufficiently robust enough to accurately segment the fiducial in all the clinical images, as shown in Figure 4 [7]. In addition to the robust segmentation, we have further made three primary improvements to improve the accuracy of C-arm tracking. (i) We have created five new designs for the fiducial with diameters varying from 5 mm – 25 mm [5]. These designs improve the segmentation accuracy not only by avoiding the intersection of the lines & ellipses in the image, but also by reducing the size of the fiducial (a smaller fiducial stays completely inside the workspace providing full feature imaging). The preliminary results indicate that the smaller fiducials offer robustness in segmentation, without any loss in the accuracy of C-arm tracking. (ii) We have proposed the use of the segmented seed locations in the X-ray images to refine the C-arm poses [4]. This information is readily available without altering the current pipeline [6]. Though the method was sufficiently robust to track the C-arms without the use of the FTRAC fiducial, the presence of the fiducial made the full algorithm practically bullet-proof. (iii) Inspired by the previous result, we investigated the possibility of using just a single ellipse in addition to the segmented seed locations to automatically track a C-arm [3]. A single ellipse is otherwise not

sufficient to track a C-arm, but in combination with the seed locations, can drive the problem home. Preliminary results indicate a tracking accuracy better than a degree. Further experiments are currently underway, the results of which would be published next year.

Figure 5: The matching problem can be converted to the flow problem, the min-cost flow giving the best match.

(c) As about 7% of the implanted seeds can be hidden in the X-ray images, hidden seeds have profound impact on the robustness of seed matching. Currently there is no feasible algorithm available that can solve this problem

robustly, without any constraints on where the images were taken from. Hence the theoretical framework of the original algorithm [11] was extended to solve the problem [8]. The min-cost flow formulation was extended for this purpose. Sets A, B, C have a node for each seed. A flow of '0' means that the two seeds do not correspond, while a flow of '1' means that they do. The links connecting S to A, (or B to C, C to A, A to T) have the capability to carry flow that could exceed 1 unit, ensuring that each seed in each image has to be matched, and it can possibly have hidden seeds at that location. The maximum flow allowed on these links yield the maximum number of seeds that can overlap. Now a min-cost flow equal to the number of seeds is pushed in the network, which achieves correct matching including hidden seeds. In essence, our novel approach exploits that although the problem is NP-hard and the network is not able to constrain the hidden seeds in any one image, still hidden

seeds in other images are extremely well constrained. Thus the min-cost flow nearly solves the problem, except for a few cases which exhibit some ambiguity. These cases usually incarnate themselves as self-consistent complete subsets of 2-5 seeds. Each of these subsets can be independently solved using an extremely fast brute-force type of algorithm. This approach can be extended to any number of images. Thus MARSHAL effectively recovers the hidden seeds and resolves the correspondence in polynomial time. Phantom experiments were conducted on a precisely fabricated acetol seed phantom, with the number of seeds ranging from 40 to 100 in increments of 15, while keeping the seed density at 1.56 seeds/cc. Four images gave a matching rate of ~100% and a worst case reconstruction error of 0.8 mm for the occasionally mismatched seeds. Thus Hidden-MARSHAL with 4 images appears clinically viable. Furthermore, MARSHAL was also extended to utilize seed orientation [9]. The original MARSHAL used only the information about seed centers, disregarding the information about the orientations and length of seeds, which could have some dosimetric impact for certain seed/isotope types and can contribute to improving seed matching rate and 3D implant reconstruction accuracy. It can also become critical information when hidden and spuriously segmented seeds need to be matched, where reliable and generic methods are not yet available. Expecting orientation information to be useful in reconstructing large and dense implants, we have developed a method which incorporates seed orientation information into MARSHAL. Simulation experiments indicated that using seed orientation also improved the

matching rate. Currently we are making further extensions to the matching algorithm by designing a tomosynthesis based algorithm that can reconstruct the seeds without explicitly segmenting them. This algorithm is currently still in development, though recent lab experiments have indicate that the algorithm can reconstruct seeds using 6 images.
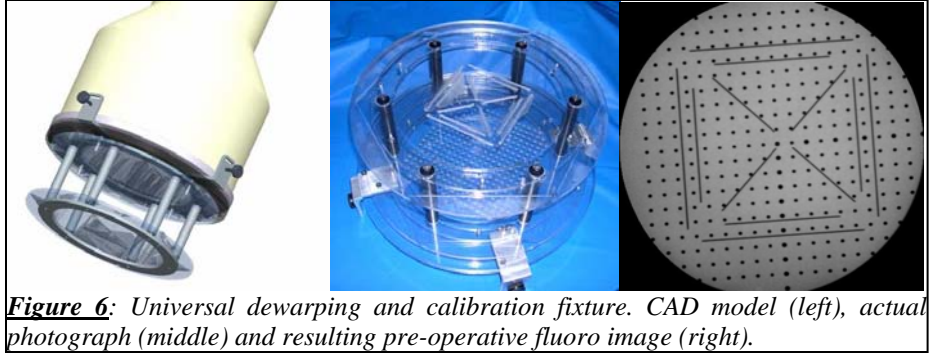
__Figure 6__: Universal dewarping and calibration fixture. CAD model (left), actual photograph (middle) and resulting pre-operative fluoro image (right).

(d) In addition to the above aspects of Aim1, we have also made further progress in the issue of C-arm calibration. Since the past submission, a new clinical-grade fully automated calibration toolkit was designed, built, and tested (Figure 6) [2]. The objective of "C-arm Calibration" is to discern intrinsic and extrinsic parameters (focal distance, etc) of the fluoroscope, achieving fully automated calibration within five minutes time, including setup and image acquisition. We developed a highly portable two-plane calibration phantom, the first plane of which is a distortion correction grid. The phantom easily fits on any C-arm with the use of a "portable flange" that allows the phantom to go on and off the C-arm within seconds. Residual errors for distortion correction showed an average value 0.26 pixels and a deviation of 0.08 pixels. Extrinsic calibration is also performed with the same phantom using the same image. A second plane on the phantom contains precision machined straight lines. Each line in space, along with its projection in the image, defines a unique 3D plane. Since the X-ray source has to pass through all these planes, the least square minimization intersection of these 10 planes provides an accurate estimate of the intrinsic C-arm parameters. Experiments using this auto-calibration phantom exhibited accuracy within 10 mm for the focal length, which is significantly better than that required for accurate 3D reconstruction of the brachytherapy sources. The total running time for the toolkit, including image capture, automatic segmentation and calculation was consistently below 1 minute. Furthermore, in a recent paper [10], we have modeled the effects of mis-calibration on 3D reconstruction as an affine transformation, and we proved its validity experimentally. We have derived bounds on the amount of scaling, translation, and rotation error. For pose-dependent calibration, we proved that using the mean calibration minimizes the reconstruction variance. Phantom experiment with the FTRAC fiducials indicated that C-arm tracking is insensitive to mis-calibration. We also proved that mis-calibration up to 50 mm adds only negligible error in 3D reconstruction of small objects, beyond which the reconstructed object start drifting with

respect to the fiducials, while retaining relative pose. Thus we have theoretically proved that a significant family of applications involving the reconstruction of small objects, such as brachytherapy implants, can function without cumbersome super-accurate on-line calibration. A constant off-line calibration should suffice.

In Aim2: we proposed to devise a procedural workflow to incorporate the RUF package with the Interplant system in a manner that does not impart significant changes on the usual flow of treatment. The proposed final clinical procedure is as outlined below. It can be observed that it does not cause any significant alteration to the current standard of care workflow and does not add in excess of 10 minutes to the current surgical time.

1. Setup:
   a. Prepare patient
   b. Assemble Interplant®
   c. Position C-arm above the patient
   d. *Attach FTRAC fiducial (NEW STEP)*
2. Insertion Planning
   a. Acquire TRUS volume of target anatomy
   b. Transfer TRUS to Interplant®
   c. Optimize insertion and dosimetry plans
3. Needle Insertion and Seed Deployment
   a. Insert needle
   b. Confirm needle location
   c. Drop seeds
   (Repeat from step 3 for a batch of needles)
4. Fluoroscopy
   a. Position C-arm
   b. Acquire and store image
   (Repeat from step 4 for each image)
5. *Registration of Ultrasound to Fluoroscopy (NEW STEP)*
   a. *Import fluoro images onto laptop*
   b. *Calibrate C-arm*
   c. *Recover C-arm poses*
   d. *Match & reconstruct seeds*
   e. *Transform seeds to template frame*
   f. *Export seeds*
6. Dosimetry & Optimization
   a. Import seed locations to Interplant®
   b. Analyze dosimetry
   c. EXIT if dose is optimal
   d. Otherwise optimize the remainder of the plan
   e. Repeat from 3a

For maximum safety and traceability, the above workflow is kept identical to that of a standard implant procedure, except adding two steps: mounting the FTRAC fiducial and performing Registration of Ultrasound to Fluoroscopy (RUF). RUF will be implemented as stand-alone software running aside the Interplant on a shared workstation. The standard implant workflow will halt after Step 4 and RUF will run separately. RUF will export a list of seeds (stored in external data file) for Interplant. Our design guarantees that (1) Interplant will not need to be modified to accommodate for RUF and (2) RUF will not jeopardize operational safety in the current Interplant. Note that the corresponding basic import function in Interplant is currently available and it is used for importing seed locations from a CT-based post implant dosimetry package. The current software implements the following: First, we remove the already implanted seed from the latest implant plan. Next, we

add the seeds to the dose field that were reconstructed from fluoroscopy. The seeds imported from fluoro reconstruction are "permanent", in the sense that they cannot be modified (because they were already implanted), contributing a permanent dose. Hence the total dose will be calculated from the union of imported and yet to be implanted seeds.

The workflow within RUF, as also illustrated in Figure 7, has been implemented as follows: We will import the recorded C-arm images from the DICOM storage. If the particular C-arm does not have its own DICOM storage, we will provide one by grabbing the video frame from the C-arm and creating a DICOM image with correct header and time stamp. We will recover the pose of the C-arm images relative to the FTRAC fiducial. We will match and reconstruct the seed locations in FTRAC space. We will transform the seed locations from FTRAC space to Interplant template space by applying the a priori known rigid transformation matrix implemented by the precision-machined extender. We will then export the seed locations from RUF to Interplant. The one-way interface between RUF and Interplant will take place through an ASCII clear text data file containing a variable list of 3D points (xyz coordinates) representing 3D seed locations. The data file will be exported from RUF automatically and imported to Interplant by operator request. For maximum flexibility and rapid prototyping ability, we have implemented RUF in the MATLAB environment. This environment has been sufficiently fast, robust, and most of all, it has been providing fast turn-around in prototyping. While MATLAB is excellent for algorithmic work, it also has a rich image processing toolbox and handles DICOM images as well.
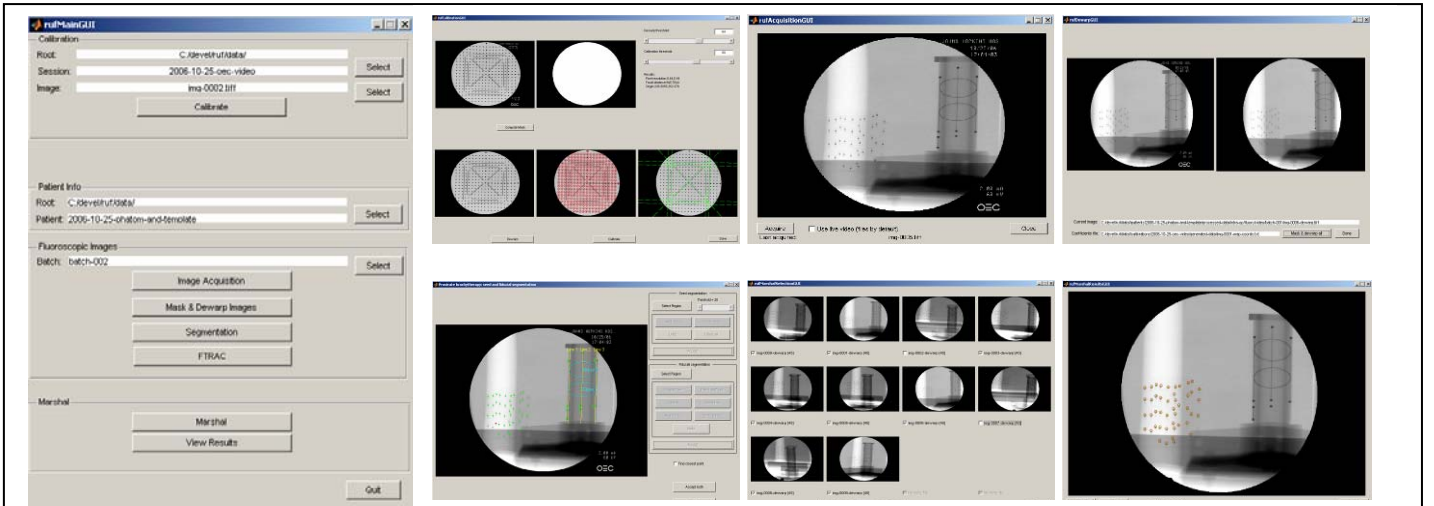
Figure 7*: GUI screen captures of the various phases of the algorithm, including GUI for the main program, calibration, image acquisition, distortion correction, seed & FTRAC segmentation, seed matching and validation.*

In Aim3: we proposed to evaluate the workspace constraints and the overall performance of the system. In order to minimize image acquisition time, while still maximizing reconstruction accuracy, the following set of optimal imaging poses have been empirically evaluated: Image-#0: 5º left lateral; Image-#1: maximum left lateral (15º); Image-#2: maximum right lateral (5º) without hitting the FTRAC; Image-#3: maximum forward tilt (10º) with 5º left lateral; Image-#4: maximum backward tilt (10º) with 5º left lateral. Though detailed experimental evaluations have been allocated in year 2, we have in conjunction with our clinical partners, conducted one dummy-surgical trial on a custom made prostate phantom (Figure 8). 45 dummy Palladium seeds were inserted using the
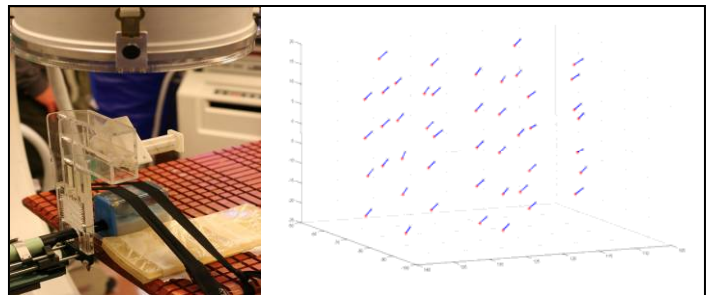
Figure 8*: The experimental setup for the procedure (left). The C-arm and Ultrasound probe are also visible. The shift between the reconstructed (red) and true seed locations (black).*

Interplant system and the full protocol carried out. To establish ground truth, the phantom was rigidly strapped to the FTRAC fiducial and a post-operative CT (0.5mm pixel size, 1 mm slices) was conducted. The seeds reconstructed from our procedure were compared to those from CT. The experiment provided an average accuracy of 1.6 mm, clinically sufficient for the procedure. It was also noticed that the reconstruction accuracy of the implant was 0. 22 mm, though the whole implant was shifted by 1.4 mm, as illustrated in the Figure 8. The possible reasons for this are currently under study, since it can potentially improve the performance.

C KEY RESEARCH ACCOMPLISHMENTS
1. Ability to intra-operatively compute 3D seed locations in prostate brachytherapy.
2. Minimal alterations to the current prostate brachytherapy clinical protocol.
3. An easy to use GUI (graphical user interface) based integrated clinical research system for validation.
4. Furthering some core technical ideas in 3D X-ray reconstruction that could be applicable in other X-ray based clinical procedures like those in interventional orthopedic and angiographic surgeries.

D KEY TRAINING ACCOMPLISHMENTS
1. Close interaction with clinicians, medical physicists, and industry advisors.
2. Various international training courses, workshops & conferences in medical imaging processing.
3. Ability to do independent research, including mentoring undergraduate students for their research.

E REPORTABLE OUTCOMES
1. Publications in leading conferences and journals, as listed in the references section (including two flagship conference publications in MICCAI and two journal publications in Medical Physics).
2. A prototype for subsequent commercialization and integration into the Interplant software.

F CONCLUSIONS
In conclusion, we have addressed the issue of intra-operative reconstruction of brachytherapy seeds, with minimal alteration to the current clinical protocol or any significant increase in cost. Preliminary experiments indicate a 3D accuracy of 1.6 mm mean and 0.2 mm standard deviation, which is acceptable for a successful clinical outcome. A prototype system for clinical trails and potential commercialization has also been tested & implemented.

So what: The success of brachytherapy chiefly depends on our ability to intra-operatively cover the prostate with sufficient radiation while still avoiding excessive radiation to surrounding organs. Currently, such level of precision is not always achievable even by the most experienced physicians. Thus many implants fail or cause severe side effects owing to faulty seed placement, a problem what still cannot be corrected in the operating room. Our results indicate the feasibility of a system that could achieve intra-operative localization of the implanted seeds in relation to the prostate, to allow for in-situ dosimetric optimization and exit dosimetry. This ability to perform intra-operative dosimetry may change the standard of care in brachytherapy by allowing the physician to achieve technically excellent brachytherapy implants, resulting in improved disease control and quality of life for a large and steadily growing group of patients.

G REFERENCES (List of publications arising from the grant, sorted by publication date)

1. [Jain-2007] **A Jain**, M An, N Chitphakdithai, G Chintalapani, G Fichtinger: *C-arm calibration - is it really necessary?* SPIE Medical Imaging 2007 (accepted), to be presented in **Feb 2007**, San Diego.

2. [Chintalapani-2007] G Chintalapani, **A Jain**, RH Taylor: *Statistical Characterization of C-arm Distortion with Application to Intra-Operative Distortion Correction.* SPIE Medical Imaging 2007 (accepted), to be presented in **Feb 2007**, San Diego.

3. [Burkhart-2007] D Burkhardt, **A Jain**, G Fichtinger: *A Cheap and Easy Method for 3D C-arm Reconstruction using Elliptic Curves*. SPIE Medical Imaging 2007 (accepted), to be presented in **Feb 2007**, San Diego.

4. [Jain-2006a] **A Jain**, G Fichtinger: *C-arm Tracking and Reconstruction without an External Tracker*. Ninth International Conference on Medical Image Computing and Computer-Assisted Intervention, Proceedings Lecture Notes in Computer Science 419, pp 494--502, Springer, **Oct 2006**

5. [Jain-2006b] **A Jain**, I Iordachita, G Fichtinger: *A small size fluoroscope tracking fiducial.* 6th Annual Meeting of the International Society for Computer Assisted Surgery, Montreal, Canada, pp 235-238, **June 2006**

6. [Vikal-2006a] S Vikal, **A Jain**, A Deguet, DY Song, G Fichtinger: *Seed segmentation in C-arm fluoroscopy for brachytherapy implant reconstruction*, Conference of American Association of Physicists in Medicine, Journal of Medical Physics, **June, 2006**. Abstract in Medical Physics, Volume 33(6), page 2229.

7. [Vikal-2006b] S Vikal, **A Jain**, A Deguet, DY Song, G Fichtinger: *Automated segmentation of radiographic fiducials for C-Arm tracking*, Conference of American Association of Physicists in Medicine, Journal of Medical Physics, **June 2006**. Abstract in Medical Physics, Volume 33(6), page 2208.

8. [Kon-2006] RC Kon, **A Jain**, G Fichtinger: *Hidden seed reconstruction from C-arm images in brachytherapy*, IEEE International Symposium on Biomedical Imaging, Arlington, VA, 2006, pp 526–29, 6-9 **April 2006**, Arlington.

9. [Zhou-2006] Y Zhou, **A Jain**, GS Chirikjian, G Fichtinger: *Incorporating Seed Orientation in Brachytherapy Implant Reconstruction*, SPIE Medical Imaging 2006:: Visualization, Image-Guided Procedures, and Display. Edited by Cleary, Kevin R.; Galloway, Robert L., Jr. Proceedings of the SPIE, Volume 6141, pp. 438-452, **Feb 2006**.

10. [Jain-2005a] **A Jain**, R Kon, Y Zhou, G Fichtinger: *C-arm calibration - is it really necessary?* Eighth International Conference on Medical Image Computing and Computer-Assisted Intervention -- MICCAI 2005, Proceedings in Lecture Notes in Computer Science Vol. 3749, pp 639-646, Springer, **Oct 2005**

11. [Jain-2005b] **A Jain**, Y Zhou, T Mustufa, EC Burdette, GS Chirikjian, G Fichtinger: *Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)*. **Nov 2005**, Med Phys, Nov; 32(11):3475-92.

12. [Jain-2005c] **A Jain**, T Mustufa, Y Zhou, EC Burdette, GS Chirikjian, G Fichtinger: *A robust fluoroscope tracking (FTRAC) fiducial*. **Oct 2005**, Med Phys. Oct; 32(10):3185-98

C-arm Calibration – is it really Necessary?

Ameet Jain, Michael An, Nicha Chitphakdithai, Gouthami Chintalapani, Gabor Fichtinger
Johns Hopkins University

Abstract

C-arm fluoroscopy is modelled as a perspective projection, the parameters of which are estimated through a calibration procedure. It has been universally accepted that precise intra-procedural calibration is a prerequisite for accurate quantitative C-arm fluoroscopy guidance. Calibration, however, significantly adds to system complexity, which is a major impediment to clinical practice. We challenge the status quo by questioning the assumption that precise intra-procedural calibration is really necessary. To derive bounds on reconstruction errors as a function of mis-calibration, we model the error using an affine transform. This is fairly intuitive, since small amounts of mis-calibration result in predictably linear transformation of the reconstruction space. Experiments indicate the validity of this approximation even for 50 mm mis-calibrations.

The primary results presented here are related to virtual fluoroscopy (VF), one of the most commonly proposed techniques for providing intra-operative fluoroscopic guidance. The problem is twofold: (a) C-arm calibration; and (b) C-arm distortion correction. Using our theoretical and experimental analysis on mis-calibrated C-arms, we propose a simple solution that could easily address the first issue, indicating that calibration is not required at all. Furthermore, to address the problem of pose dependant distortion correction, we propose a framework that can statistically study the maximum variation in distortion near a certain pose and then intra-operatively use an average correction. These theoretical derivations and experimental results make a strong case for the use of mis-calibrated C-arms, obviating the cumbersome intra-operative calibration process, potentially boosting clinical applicability of quantitative fluoroscopy in many procedures.

Keywords: Calibration, C-arms, fluoroscopy guided procedures
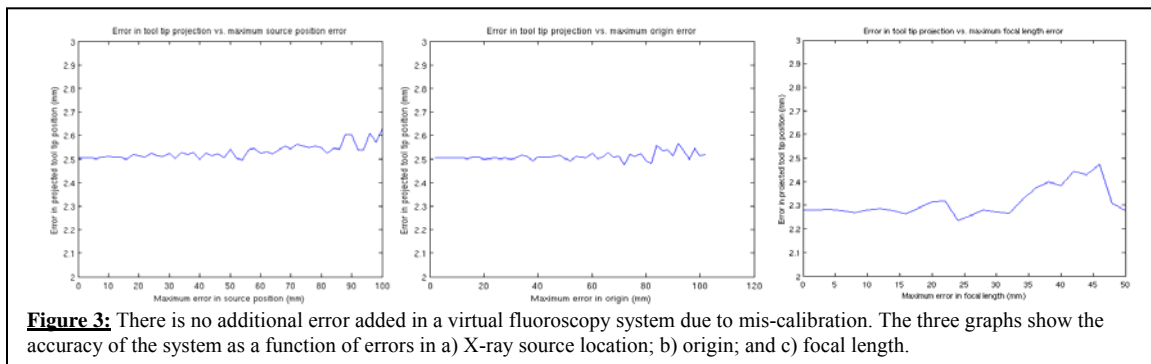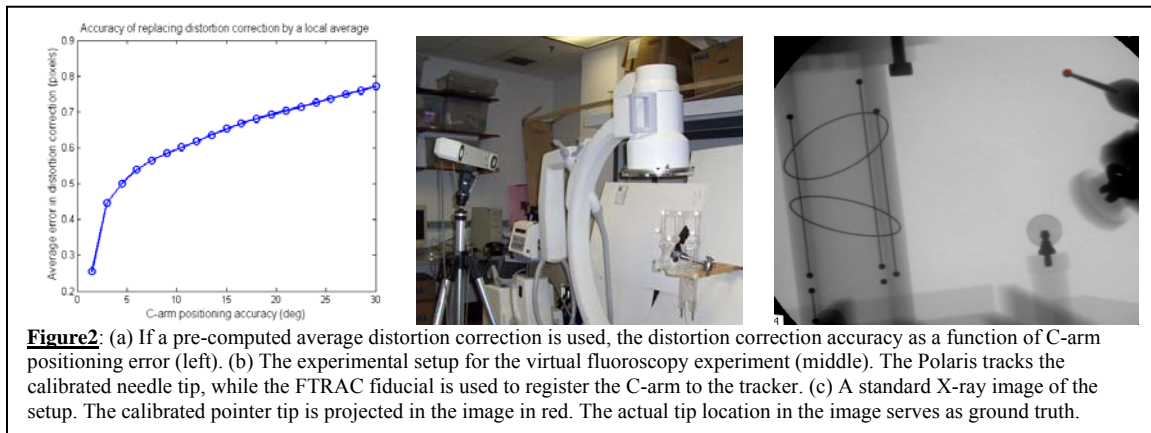
Description of purpose: Quantitative C-arm fluoroscopy measures the (relative) spatial transformation between two objects, as compared to the transformation between an object and the C-arm. Thus the central intuition is that *while an incorrect calibration gives erroneous estimates for the absolute transformations, nevertheless it still provides acceptable relative estimates*, implying that intra-operative calibration may not be needed.

Figure1: Mis-calibration shifts all reconstructed objects under an affine transformation. Relative location of objects is preserved.

Method: To address the problem of pose dependant distortion correction (DC), data should be collected *'near'* a certain pose pre-operatively. A principal component analysis on this data provides useful statistical information like the mean, primary deformation modes, etc. A simple strategy would be to use just the mean DC function *'near'* a certain pose (most applications have a pre-determined protocol for taking X-ray images). To address intrinsic camera calibration, note that only relative transformations matter. Hence we propose that the C-arm to tracker registration should be done using a fiducial object near the patient region of interest (the full theoretical proof for this observation would be provided in the full paper).
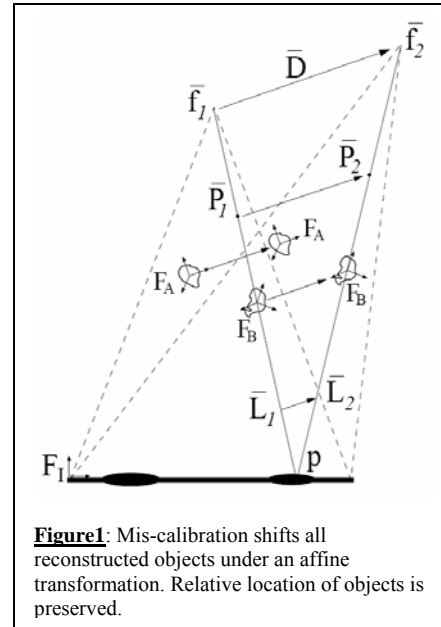
Experiments and Results: X-ray images of a calibration phantom with 300 BBs were taken in a C-arm sweep every three degrees. Near any desired C-arm position, a PCA is performed on the DC function. Three modes were found to capture 99% variation. The mean was used to model distortion at every pose within a cone near the desired location and the error noted.

A VF setup was installed using the Polaris as the tracker and an accurate C-arm calibration phantom. The fluoroscope tracking (FTRAC) fiducial (0.6 mm translation, 0.3° rotation accuracy) was used to register the C-arm to the Polaris. Randomly chosen known calibration errors were added and the error of the VF guidance system plotted as a function of these errors.

Figure2: (a) If a pre-computed average distortion correction is used, the distortion correction accuracy as a function of C-arm positioning error (left). (b) The experimental setup for the virtual fluoroscopy experiment (middle). The Polaris tracks the calibrated needle tip, while the FTRAC fiducial is used to register the C-arm to the tracker. (c) A standard X-ray image of the setup. The calibrated pointer tip is projected in the image in red. The actual tip location in the image serves as ground truth.

Figure 3: There is no additional error added in a virtual fluoroscopy system due to mis-calibration. The three graphs show the accuracy of the system as a function of errors in a) X-ray source location; b) origin; and c) focal length.

Conclusions: Fig 2(a) indicates that with a 10° C-arm re-positioning accuracy, using a mean distortion correction function leads to only a 0.6 pixel (0.2 mm) average distortion error at each pixel. This suggests that in many applications, a mean DC function 'near' a certain pose would suffice. From Fig. 3, we can conclude that mis-calibration errors as high as 50-100 mm do not seem to add any additional error in a virtual fluoroscopy setup. This suggests that using our suggested protocol, intra-operative C-arm calibration might not be necessary.

New or breakthrough work to be presented: A theoretical framework along with experimental results to prove that intra-operative C-arm calibration is not required in a host of surgical applications.

Indicate whether the work is being, or has been, submitted for publication or presentation elsewhere, and, if so, indicate how the submissions differ: Some preliminary ideas have been published last year, but the full framework, including these *new* results for VF have not been submitted anywhere.

Statistical Characterization of C-arm Distortion with Application to Intra-Operative Distortion Correction

Gouthami Chintalapani, Ameet K. Jain, Russell H. Taylor
Department of Computer Science, Johns Hopkins University

C-arm images undergo significant pose dependant distortion, which needs to be corrected for intra-operative quantitative 3D surgical guidance. Several distortion correction techniques have been proposed in the literature, the current state of art uses a dense grid pattern rigidly attached to the detector. These methods become cumbersome for intra-operative use, such as 3D reconstruction, since the grid pattern interferes with patient anatomy. The primary contribution of this paper is a framework to statistically analyze the distortion pattern which leads to alternate intra-operative distortion correction methods. In particular, we propose a new phantom that uses very few BBs, and yet accurately corrects for the distortion.

The high dimensional space of distortion pattern can be effectively characterized by principal component analysis (PCA), with only the first three eigen modes capturing about 99% of the variation. Phantom experiments indicate that distortion map can be recovered up to an average accuracy of about 0.1mm/pixel with these three modes. With this prior statistical knowledge, a subset of BBs can be sufficient to recover the distortion map accurately. Simulations indicate that 14 BBs can recover distortion with average error of 0.14mm/pixel, accuracy sufficient for most clinical applications. These BBs can be arranged on the periphery of the C-arm detector, minimizing the interference with patient anatomy and hence allowing the grid to remain attached to the detector during surgery. The proposed method is fast, economical, convenient to use intra-operatively and independent of the C-arm, potentially boosting the clinical viability of applications such as quantitative 3D fluoroscopic reconstruction.

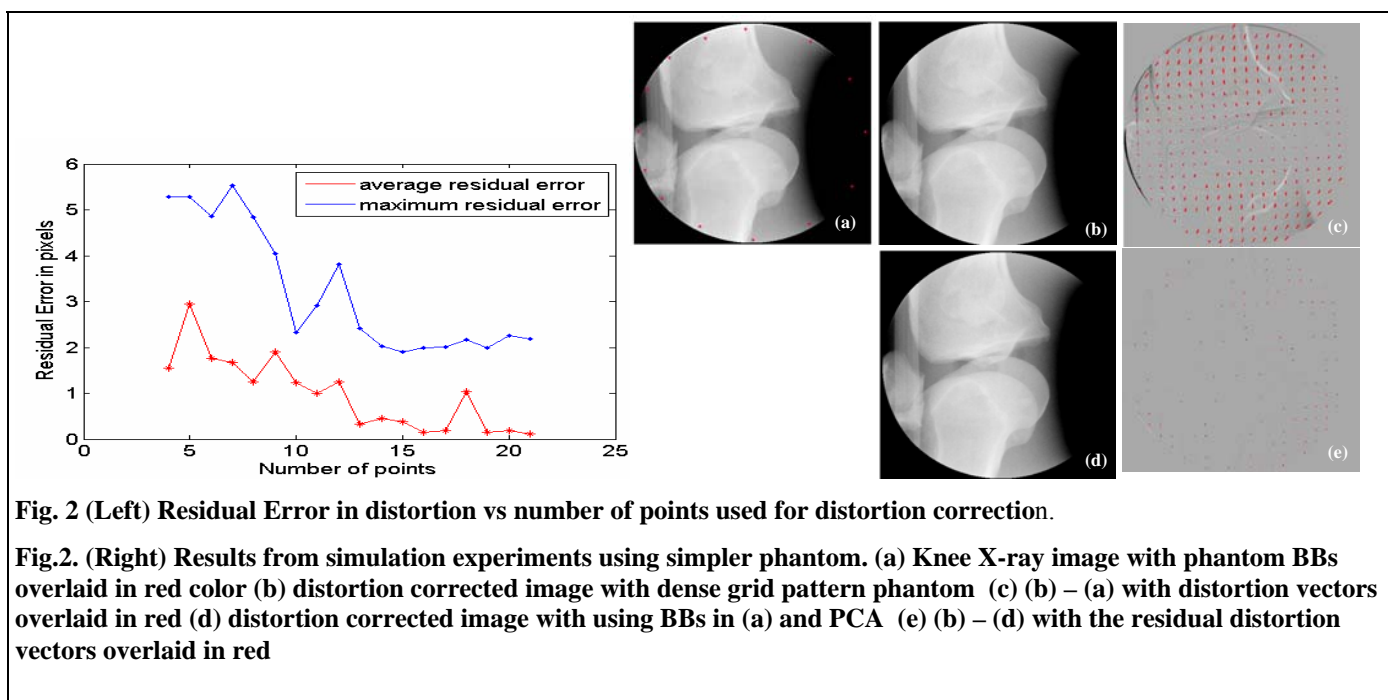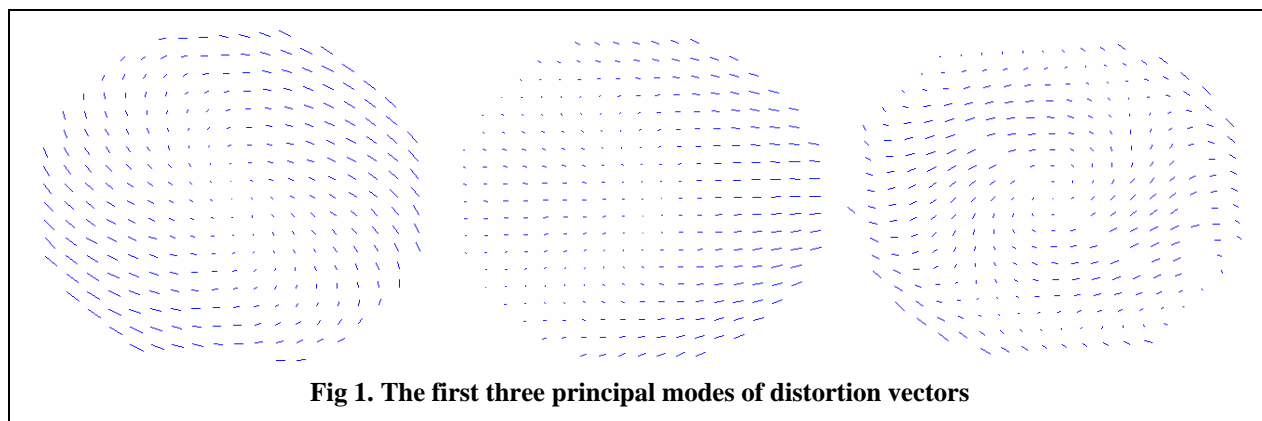*Keywords: Methods*: Quantitative image analysis; *Modalities*: X-ray CT; *Other*: distortion correction

Presentation preference: Short Talk (poster presentation acceptable)

Description of purpose: C-arm fluoroscopy is the most commonly used intra-operative imaging modality because of its low cost and portability. C-arm images suffer from distortion and the distortion varies from place to place and from pose to pose. The most common distortion correction method is to use a grid pattern BBs or lines on a plate attached to the detector. These methods become cumbersome for intra-operative use (such as for 3D reconstruction), since the grid pattern either significantly corrupts the patient's image, or necessitates two X-ray shots at each pose (one with and the other without the grid). In this paper, we propose a statistical framework that enables us to do intra-operative distortion correction of the C-arm X-ray images using a simpler phantom with fewer BBs.

Methods: X-ray images of the grid-pattern phantom with 300 BBs were taken in a C-arm sweep along the propeller axis, one every three degrees. Distortion was modeled with a fifth order 2D Bernstein polynomial with 72 vector coefficients. PCA is performed on the distortion vectors associated with these images. The first three principal modes capture 99% of the variation and are shown in Fig 1

Experiments and Results: Simulation experiments were carried out on the images acquired from OEC 9600 C-arm. PCA is performed on 104 images with 10 images used in leave-out tests. The distortion parameters can be recovered up to an average accuracy of 0.1mm/pixel and maximum error is about 0.3mm. Validation experiments on the number of BBs required show that eight BBs are sufficient to do distortion correction with an average error of 0.2mm per pixel. These results show that distortion patterns, although complicated, are predictable with PCA.

Fig 1. The first three principal modes of distortion vectors

Fig. 2 (Left) Residual Error in distortion vs number of points used for distortion correction.

Fig.2. (Right) Results from simulation experiments using simpler phantom. (a) Knee X-ray image with phantom BBs overlaid in red color (b) distortion corrected image with dense grid pattern phantom (c) (b) – (a) with distortion vectors overlaid in red (d) distortion corrected image with using BBs in (a) and PCA (e) (b) – (d) with the residual distortion vectors overlaid in red

New or breakthrough work to be presented: A very high dimensional C-arm distortion is statistically explained with the first three principal components. With this prior statistical knowledge, a significantly simpler phantom can be used to do intra-operative distortion correction.

Conclusions: A statistical framework to characterize complex distortion patterns is presented. We extended this framework to develop a new simple phantom based intra-operative distortion correction technique. More experiments are being conducted to validate this framework like number of images needed to extract principal modes, to study if and how the principal modes change over time or across C-arms etc. and the results will be included in the full paper.

Indicate whether the work is being, or has been, submitted for publication or presentation elsewhere, and, if so, indicate how the submissions differ: This work has never been published before and is being submitted for the first time to SPIE Medical Imaging

A Cheap and Easy Method for 3D C-arm Reconstruction using Elliptic Curves

David Burkhardt[1], Ameet Jain[2], Gabor Fichtinger[2]

1- Haverford College,   2 - Johns Hopkins University

Abstract

For quantitative C-arm fluoroscopy, we had earlier proposed a unified mathematical framework to tackle the issues of pose estimation, correspondence and reconstruction, without the use of external trackers. The method used randomly distributed unknown points in the imaging volume, either naturally present or induced by randomly sticking beads on the patient. These points were then inputted to an algorithm that computed the 3D reconstruction. The algorithm had an 8^o region of convergence, which in general could be considered sufficient for most applications. Here, we extend the earlier algorithm to make it more robust and clinically acceptable.

We propose the use of a circle/ellipse, naturally found in many images (ex: as the head of an inserted screw). Even if induced artificially, such curves are extremely easy and cheap to manufacture. We show that the projection of elliptic curves (also elliptic curves) constrain 5 out of the 6 degrees of freedom of the C-arm pose. To completely recover the true C-arm pose, we use an additional constraint in the form of any available point correspondences in the images. We provide an algorithm to easily obtain a virtual correspondence across all the images and show that two correspondences can always recover the true pose uniquely. Preliminary phantom experiments across three images indicate a pose estimation accuracy of 1.9^o and 3D reconstruction accuracy of 1.5 mm using only an ellipse and two point correspondences. The method appears to be sufficiently accurate for many clinical applications and particularly appealing since it works without any significant modification of the clinical protocol.

<u>Keywords</u>: Tracking, Localization, Registration, X-ray reconstruction, C-arm pose estimation

Description of purpose: Quantitative fluoroscopy has not found a *large scale clinical acceptance*, because of inherent technical difficulties involving intra-operative calibration of model parameters, pose tracking, and target matching/reconstruction. Advanced commercial and academic systems employ resident calibration structures and optical/electromagnetic trackers or calibrated radiographic fiducials to obtain the C-arm pose. The resulting equipage tends to be prohibitively expensive and complex that often interferes with the subject, image space, and clinical work-volume. In many applications, screw/needle ends, implanted surgical markers, special anatomy points etc. are naturally present in the images. By enforcing the "consistency" of these feature points across the images, one can potentially solve for all unknown parameters of calibration, pose recovery, matching, and reconstruction in one massive high-dimensional non-linear optimization. We propose such a framework in this paper, in particular, where an elliptic curve is visible in the images.

Method: We model the complete 3D reconstruction problem as a mathematical optimization problem, as shown in Fig. 1. An image of an elliptic curve can be used to infer a 3D cone in

space on which the elliptic curve must lie. Through a series of mathematical transformations on this 3D cone that incorporate known physical properties of the elliptic curve, we constrain five of the six degrees of freedom required for pose reconstruction. An optimization algorithm can now use any additional constraint to lock into the correct pose. The constraint that we use is a correspondence across the images of anatomical points.

$$\begin{array}{c} \arg\min \\ \mathcal{R}, \mathcal{T}, \mathcal{M}, f \end{array} \quad \sum_{i=1}^{N_1}\sum_{j=1}^{N_2}\sum_{k=1}^{N_3} \mathcal{C}_{ijk} f_{ijk}$$
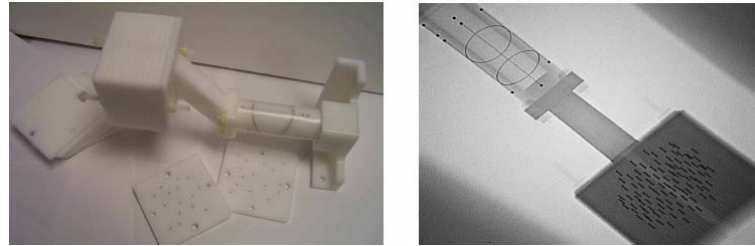
$$\text{where} \quad \sum_{i=1}^{N_1}\sum_{j=1}^{N_2} f_{ijk} \geq 1 \; \forall \; k \; ; \quad \sum_{j=1}^{N_2}\sum_{k=1}^{N_3} f_{ijk} \geq 1 \; \forall \; i;$$

$$\sum_{k=1}^{N_3}\sum_{i=1}^{N_1} f_{ijk} \geq 1 \; \forall \; j \quad and \quad f_{ijk} \in \{0,1\}$$

Figure1: The full mathematical description of the 3D reconstruction process. The best reconstruction and pose parameters would minimize the back-projection errors.

Experiments and Results: A radiographic fiducial was used to track the C-arm (0.56 *mm* translation; 0.33^o rotation accuracy), and was accurately attached to a point cloud phantom as shown in Fig 2. Images of this phantom were taken 25^o around the AP-axis using a *Philips Integris V3000* fluoroscope. Each image was hand segmented to establish the true segmentation and correspondence. This was used to create the ground truth. The true C-arm pose and reconstruction was

Figure3: The high-precision experimental phantom (left). 3D locations of any feature wrt any other feature are known precisely by accurate fabrication. A sample X-ray image (right) of the phantom

compared to that computed using the proposed algorithm. A sample set of computed poses is presented in Fig. 3. Preliminary results from these phantom experiments indicate that we can on
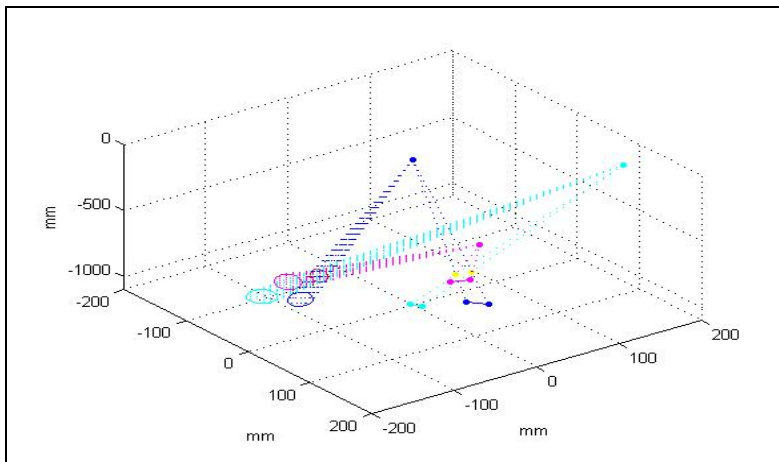
Figure 3: Three sample C-arm images (right) and the relative C-arm poses reconstructed from them using a single ellipse and two point correspondences (above). Cones vertices represent x-ray sources, while cone bases and projected points lie in the image planes.

average track a C-arm to a rotational accuracy of 1.9^o using only two point correspondences. The algorithm achieves a 3D reconstruction error of 1.5 mm, an accuracy sufficient for most clinical applications. Furthermore, additional correspondences will exist in most clinical applications; these correspondences further decrease this error. The algorithm has a large region of convergence and is free of initial estimates.

Conclusions: A framework using elliptic curves for C-arm tracking and reconstruction without compromising on the available work-volume has been proposed. Experiments indicate an accuracy and robustness sufficient for most clinical applications. The framework obviates expensive external tracking devices and is capable of using information naturally present in a family of procedures. In applications where this information is not present, the framework easily extends by randomly attaching an ellipse or a few beads near the patient anatomy.

New or breakthrough work to be presented: We propose a novel method to use a randomly placed elliptic curve to automatically track a C-arm and reconstruct in 3D any anatomy of interest. This method is both more economically and clinically feasible, when compared to most contemporary techniques.

Indicate whether the work is being, or has been, submitted for publication or presentation elsewhere, and, if so, indicate how the submissions differ: This work has not been submitted for presentation anywhere else.

C-arm Tracking and Reconstruction Without an External Tracker⋆

Ameet Jain and Gabor Fichtinger

Department of Computer Science, Johns Hopkins University

Abstract. For quantitative C-arm fluoroscopy, we have developed a unified mathematical framework to tackle the issues of intra-operative calibration, pose estimation, correspondence and reconstruction, without the use of optical/electromagnetic trackers or precision-made fiducial fixtures. Our method uses randomly distributed unknown points in the imaging volume, either naturally present or induced by randomly sticking beads or other simple markers in the image pace. After these points are segmented, a high dimensional non-linear optimization computes all unknown parameters for calibration, C-arm pose, correspondence and reconstruction. Preliminary phantom experiments indicate an average C-arm tracking accuracy of 0.9^o and a 3D reconstruction error of $0.8\ mm$, with an 8^o region of convergence for both the AP and lateral axes. The method appears to be sufficiently accurate for many clinical applications, and appealing since it works without any external instrumentation and does not interfere with the workspace.

1 Introduction

C-arm fluoroscopy is ubiquitous in general surgery, interventional radiology, and brachytherapy, due to its real-time nature, versatility, and low cost. At the same time, quantitative fluoroscopy has not found a *large scale clinical acceptance*, because of inherent technical difficulties involving intra-operative calibration of model parameters, pose tracking, and target matching/reconstruction. While these aspects have been studied extensively, a clinically extant solution appears to be lacking. Advanced commercial and academic systems employ resident calibration structures [1,2,3] and optical/electromagnetic trackers or calibrated radiographic fiducials [4,5,6] to obtain the C-arm pose. The resulting equipage tends to be prohibitively expensive and complex that often interferes with the subject, image space, and clinical work-volume. While some procedures may be more tolerant to these shortcomings, despite pressing clinical needs, quantitative fluoroscopy is completely missing from brachytherapy, which is the motivating application of our project.

It can be observed that point correspondence across images (without any knowledge about their 3D locations) is a very strong constraint for pose estimation, also referred to as *bundle adjustment* in computer vision [7]. In fact,

---

six *known correspondences* across two X-ray images are sufficient to constrain the relative C-arm pose. If eight or more correspondences are available, well known linear methods exploiting the *fundamental matrix* can be applied [7], while with five correspondences, a maximum of ten degenerate solutions are possible. Thus, in general, five point correspondences across three X-ray images can recover the C-arm locations. Theoretically four correspondences have been shown to be sufficient to recover the poses in general, barring a zero measure set of the configurations that can lead to multiple solutions [8]. Though, these configurations are known to lie on certain special cubic curves (incld. cases when three points project collinearly on any X-ray image), an intuitive understanding of all these point constellations is not yet known. Moreover, it should be noted that certain special point constellations can always be created, such that they can never be resolved uniquely using any fixed number of images.

In many applications, radioactive seeds, screw/needle ends, implanted surgical markers, special anatomy points etc. are naturally present in the images. By enforcing the "consistency" of these feature points across the images, one can potentially solve for all unknown parameters of calibration, pose recovery, matching, and reconstruction in one fell swoop, in one massive high-dimensional non-linear optimization loop. In applications that do not have an adequate information in their images, one can place a few additional sticky beads or wire markers *randomly* on any temporarily stationary part (for example on the patient skin or under the operating table), and then apply the same framework. Thus the subject of this paper is a unified mathematical framework to solve the problems of intra-operative calibration, pose tracking, and target matching, and reconstruction without any sort of pre-fabricated external fiducials or tracking instrumentation. In applications where there is no need for real-time tracking of mobile surgical instruments, for example in prostate brachytherapy, the complete elimination of intra-operative tracking and calibration entourage promises to lead to a wider clinical acceptance of quantitative 3D C-arm fluoroscopy.

2   Methods and Materials

The three integral components of this problem, in decreasing order of complexity are: (1) point correspondences; (2) C-arm pose; and (3) C-arm geometry calibration. We assume that the points have been segmented from the X-ray images. Though generic and extendible to any number of images, we currently develop the framework for exactly three images. The reason for not using two images is that they have reconstruction singularities, solved by introducing a third image, which in turn makes the problem NP-Hard (*i.e.* no algorithm can even verify the optimality of a given solution in polynomial time). The theoretical complexity of four or more images is similar to that of three images. Thus we propose a detailed solution for three images, which is easily extendible to multiple images.

2.1   Mathematical Framework

Let N be the number of points chosen arbitrarily from the clinical work volume and let N_m be the number of points detected in images I_m with pose $[R_m, T_m]$ and projection model M_m. We do not assume that the 3D points are distinctly visible in all the images, but are allowed to be hidden under one another. Though this makes the correspondence problem significantly harder, it is a more realistic representation of the clinical setting. Let s_{lm} be the position of l^{th} point in m^{th} image. When three images are used, the problem can be formulated as a large optimization problem.

$$\underset{\mathcal{R}, \mathcal{T}, \mathcal{M}, f}{\arg\min} \quad \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \sum_{k=1}^{N_3} \mathcal{C}_{ijk} f_{ijk}$$

where
$$\sum_{i=1}^{N_1} \sum_{j=1}^{N_2} f_{ijk} \geq 1 \ \forall \ k \ ; \quad \sum_{j=1}^{N_2} \sum_{k=1}^{N_3} f_{ijk} \geq 1 \ \forall \ i; \tag{1}$$

$$\sum_{k=1}^{N_3} \sum_{i=1}^{N_1} f_{ijk} \geq 1 \ \forall \ j \quad and \quad f_{ijk} \in \{0, 1\}$$

$\mathcal{C}_{ijk}$ is the the cost (described later) of matching point s_{i1} to points s_{j2} and s_{k3}. Note that it varies with any variation in $\mathcal{R}, \mathcal{T}, \mathcal{M}$. f_{ijk} is a discrete variable taking a value 1/0, and deciding the correctness of the match $< i, j, k >$. The inequalities force every segmented point to be chosen atleast once. Thus, f represents any feasible global match (and vice versa), with the cost of that correspondence given by $\sum \sum \sum \mathcal{C}_{ijk} f_{ijk}$. The problem hence is to compute $\mathcal{R}, \mathcal{T}, \mathcal{M}, f$ that minimize the total cost. It should be noted that since the images represent a real situation, this optimization has a solution with a near-zero cost. The only case in which a unique solution might not exist is when the information is not sufficient, *i.e.* when the number of beads are less than 7 or when they lie in a degenerate configuration.

Complexity: This is a non-linear optimization in $N^3 + 20$ variables with $3N^2$ constraints. The pose and model parameter optimization is in a continuous 20 dimensional space (2×6 for each pose, 3×3 for each model, one less for scale), while that for the correspondences is in a discrete combinatorial space. Note that we assume pixel sizes to be constant and known. Even if the pose & model parameters are known, it can be shown that the combinatorial optimization reduces to the minimum-weight tri-partite matching problem, known to be NP-Hard. This severely ill-conditions the problem, necessitating methods to constrain the problem adequately. It should be noted that though the global optima for two image matching can be proved to have only a cubic complexity, in many cases it suffers from singularities (Figure 2 (a)), forcing the use of a third image and hence an exponential complexity.

3D Reconstruction: Though crucial, 3D reconstruction is not explicitly incorporated into the framework. By optimizing for the cost, we also indirectly

compute the best reconstruction. Thus when the true pose parameters have been computed, the correct 3D reconstruction will come out as a byproduct. In the meanwhile, we shall only concern ourselves with the explicitly incorporated pose, correspondence and calibration parameters.

C-arm Imaging Model: C-arm imaging is typically approximated as a 5-parameter projection camera [1,2,3,4] to be calibrated intra-operatively for *each* individual image. Recently, however, it has been shown [9] that C-arm calibration might not always be necessary (Figure 1), implying that any reasonable calibration can be *assumed* for each image without actually calibrating at any time. Note that we assume the imager pixel sizes to be fixed and known.
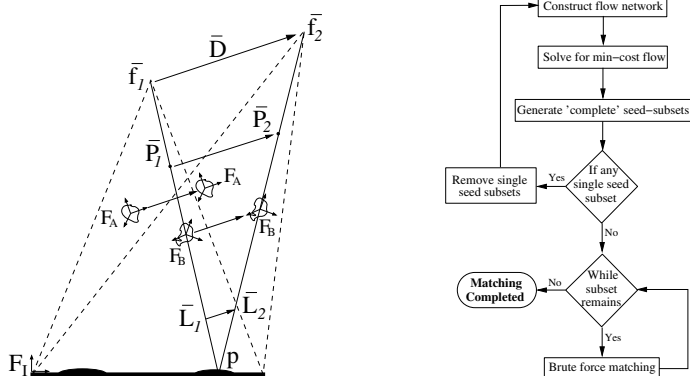

Fig. 1. Mis-calibration shifts all reconstructed objects, but keeps the relative pose nearly constant (left). The flowchart for the correspondence algorithm (right).

The central intuition is that *while an incorrect calibration gives erroneous estimates for the absolute transformations, nevertheless it still provides acceptable relative estimates.* Experimental results corroborate the theoretically derived bounds in that mis-calibration by as much as 50 mm still allows for tracking with an accuracy of 0.5 mm in translation and 0.5^o in rotation, and such mis-calibration does not impose any additional error on the reconstruction of small objects [9]. Thus to condition the optimization in Equation (1) better, it it advisable not to solve for the imaging parameters using sparse data from the image, but just use nominal values that may be known from a pre-calibration or the manual/header. An alternate perspective is to notice that since reconstruction errors change only negligibly with calibration errors, any attempt to calibrate using a sparse point set (instead of a very accurate calibration fixture) will ill-condition the problem by allowing for a whole space of feasible solutions. Thus it is wiser to fix the value at a choice that is practically close to reality. Later, if needed, the calibration can be further refined after the optimization converges.

Though we do not explicitly address the issue of distortion correction, advancements in intensifier tubes allow for lesser distortion and more recently the

advent of flat panel detectors obviates this step altogether. Furthermore, many application like brachytherapy, with their limited C-arm workspace, allow for a constant pre-operative distortion correction.

Correspondence: Assuming known pose parameters, we converted the point correspondence problem to a weighted tri-partite matching problem in Equation (1), an NP-Hard combinatorial optimization problem. An attractive approximate solution using a network-flow-based combinatorial optimization has recently been extended to efficiently deal with "hidden seeds", (i.e. points that overlap in some images) in practically $O(N^3)$ times [10].
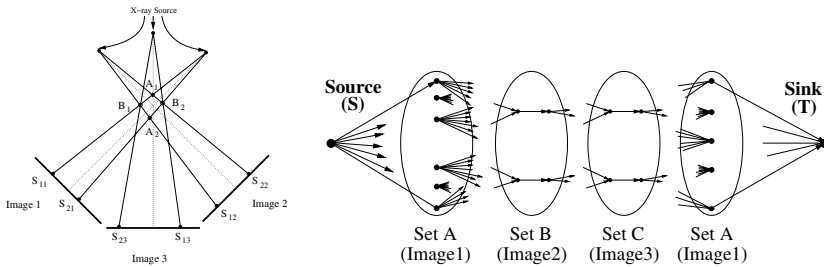
Fig. 2. A third image is needed to resolve two-image singularities (left). The flow network formulation used to solve the correspondence problem.

Sets A, B, C, and D, represent an image each. Links with a cost connect every feasible match between any two images. A flow of value N originates at the source S and ends at sink T. The problem reduces to computing a min-cost flow, easily computed using the cycle cancelling algorithm (pushes negative cycles until there are none left). A flowchart is illustrated in Figure 1. Since the problem is NP-hard the network cannot completely constrain the same point in both set A and D. Nevertheless, it works well, producing near perfect matchings.

Cost Metric: In general, any cost-metric that directly measures the deviation from the observation should perform well. The metric should incorporate all the available information, making the global minimum sharper and the algorithm robust. One good choice for a metric is *projection error* (PE). For any given set of poses and correspondence, the intersection of the three lines that join each projection to its respective x-ray source can be computed using a closed form solution that minimizes the L_2 norm of the error. PE can be computed by projecting this 3D point in each image and then measuring the distance between the projected location and the observed location of the point.

2.2   Optimization Strategy

Due to the convolution of both continuous and discrete parameters, the optimization in Equation (1) becomes ill-conditioned. Incorrect pose estimates will

invoke incorrect correspondences, especially for *dense* point clouds. However, it should be observed that *incorrect pose estimates and the subsequent correspondences invoked by it, are typically inconsistent.* This is because, while it is true that any given pose estimates invoke a correspondence, it is also true that any given correspondence also invokes a unique pose where Equation (1) will be minimized for that correspondence. Thus the desired minima will be such that the current pose invokes a correspondence, and the correspondence in turn will also invoke the same pose (with near zero cost). This order of stability we believe, will exist only at the true unique global minima. We propose an iterative strategy that exploits this observation, in spirit similar to a coordinate descent method.

Another observation to make is that given any generic estimate of the pose, the correspondence is usually completely incorrect. Nevertheless, some other pose in the vicinity can usually establish atleast a few correct correspondences ($\sim 10\%$). This new pose estimate will *behave* like a local minima. We say behave because, for any fixed pose estimate, computing f such that Equation (1) is minimized is not a polynomial time computation ($N!^2$). Thus only a working algorithm can be practically available. Even though our correspondence algorithm has been experimentally shown to be over 98% accurate near the correct pose, its assumptions start breaking down at incorrect estimates. Nevertheless, if we can estimate these few correctly matched points and block the rest, we can quickly converge to the correct answer. In the absence of any additional information, a working strategy is that if a flow in the network originates at a vertex i in set A and also ends up at vertex i in set D, then this flow is *self-consistent.* We choose a subset of self-consistent points as matched points (typically the ones closer to the average PE). These few points can now be easily used to update the pose estimates, which in turn could provide a improvement in the correspondences at a later stage. Thus the algorithm iteratively establishes the best possible correspondences (keeping the pose relatively constant) and then uses a *self-consistent* subset of points to refine the pose. As the iteration progresses, it can stop only at a self-consistent parameter choice where the pose and correspondence perfectly complement each other.

For very high density clouds, the optimization might require non-practical number of iterations to converge. Two main methods to constraint such cases are: (1) establishing a good initial estimate using prior knowledge about the surgical protocol and workspace constraints. (2) in the absence of good initial estimates, a couple of known correspondences can prove to be sufficient. These might be naturally available or artificially induced. In any case, this might become a necessary step since projective geometry can recover the 3D reconstruction only up to an arbitrary scale. To recover the scale, information external to the image is required (ex. length of an inserted screw), allowing for a few known correspondences.

3 Phantom Experiments and Results

A radiographic fiducial was used to track the C-arm (0.56 mm translation; 0.33^o rotation accuracy), and was accurately attached to a point cloud phantom as

shown in Figure 3. The cloud phantom comprises of multiple slabs, thus capable of multiple *random* point configurations. 100 points with 1.56 *points/cc* were used. X-ray images within a 20^o cone around the AP-axis were *randomly* taken using an *Philips Integris V3000* fluoroscope and dewarped. Thus both the seed locations and X-ray pose were not biased/optimized in any way, closely representing an uncontrolled surgical scenario. Each image was hand segmented to establish the true segmentation and correspondence. The true C-arm pose and reconstruction was compared to that computed from the algorithm.
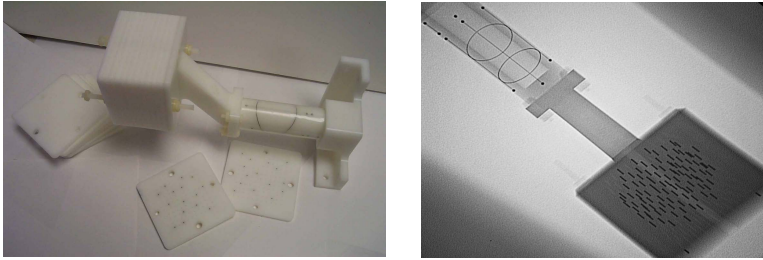

Fig. 3. An image of the point phantom attached to the fiducial (left). The phantom can replicate many point configurations. A typical X-ray image of the combination (right).

We divided the experiments into three separate cloud sizes: sparse having an average 3D point separation $\geq 25\ mm$, medium having $\sim 15\ mm$, while a high density one being $\leq 10\ mm$. These represent different types of surgical scenarios, ranging from orthopedic to brachytherapy. We generated random clouds using $10 - 20, 20 - 40\ \&\ 40 - 100$ points. Region of convergence (ROC), accuracy of C-arm tracking and reconstruction error (RE) are the three metrics used to evaluate performance. Since the scale is not directly recoverable, only the rotation errors are used to study errors in the pose. To study RE, the scale is established using two points from the fiducial.

Figure 4 plots the performance of the algorithm. Each data point is averaged using 10 random runs of the initial estimate and the point cloud. When no prior correspondences are available, the algorithm could have some difficulty in converging reliably. Pose recovery accuracies vary with the number of available points, the average being about 0.9^o, while RE remains fairly stable at $0.8\ mm$. The ROC for sparse and medium sized clouds is $8 - 10^o$ (individually along both AP and lateral axes), while it is about 6^o for dense implants.

Runtime: The algorithm was implemented in Matlab 7 on a Windows PC (3.2 GHz P4, 1GB RAM). The algorithm would typically converge in anywhere between 3-7 total iterations, taking 2-7 minutes, depending on the point cloud density and initial estimate. However, it should be noted that it spends about 50-70% of the time for file I/O (a Matlab constraint). Thus a C/C++ implementation is expected to run in 30s.
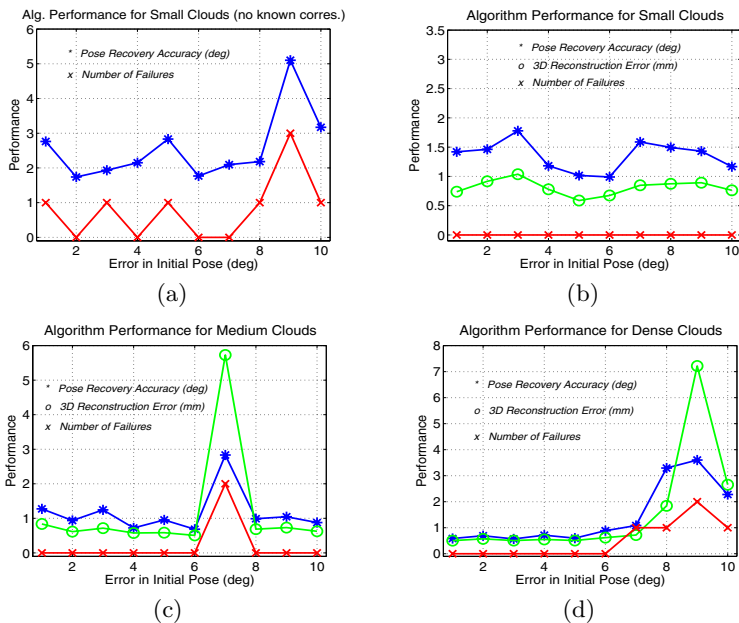
Fig. 4. The performance of the algorithm as a function of initial estimate. (a) is with no known correspondences, while (b)-(d) are with a few known ones.

4   Conclusion

A unified framework for point correspondence, C-arm tracking and reconstruction has been proposed and experimentally validated on phantoms. The experiments indicate an accuracy of 0.9^o for tracking, $0.8\ mm$ for 3D reconstruction and a convergence region of 8^o (each) in both the AP and lateral axis of rotation. The framework does not need external fiducials for C-arm pose estimation and is capable of using information naturally present in the X-ray images of a family of clinical applications, such as prostate brachytherapy. In applications where this information is not present, or a greater accuracy is desired, the framework easily extends by randomly attaching beads around the patient. Our technique does not compromise on the available clinical work volume. The framework can also accommodate any available prior information on projection angles or correspondences to constrain the optimization better, and thereby to achieve a higher accuracy.

The main concern for the clinical use of methods relying heavily on high dimensional optimization is that of providing uniformity and reliability in performance. We have conducted our validation on randomly selected views and number/distribution of the points, indicating the robustness of the algorithm to these issues. Nevertheless, this is only a first step and further work to achieve better uniformity in the results is desirable. Though the alternative of well designed

calibration fixtures and image acquisition procedures are also available, they become cumbersome in many procedures. Further development of the current approach or even an amalgamation of the two approaches, could significantly improve the current clinical viability of intra-operative quantitative fluoroscopy. The development of a clinical prostate brachytherapy system to further validate our approach is currently underway. Note that, even though the driving application was prostate brachytherapy, the method also has potential in many synergistic applications in orthopedics and angiography.

References

1. Navab, N., et al: Dynamic geometrical calibration for 3-d cerebral angiography. In: SPIE Medical Imaging. (1996) 361 – 70
2. Tang, T.: Calibration and point-based registration of fluoroscope images. Master's thesis, Queen's University (1999)
3. Livyatan, H., Yaniv, Z., Joskowicz, L.: Robust automatic c-arm calibration for fluoroscopy-based navigation: A practical approach. In: MICCAI. (2002) 60–68
4. Hofstetter, R., Slomczykowski, M., Sati, M., Nolte, L.: Fluoroscopy as an imaging means for computer-assisted surgical navigation. CAS **4(2)** (1999) 65–76
5. OEC 9800 FluoroTrak^{TM}: (GE Healthcare, Waukesha, WI)
6. Jain, A., Mustufa, T., Zhou, Y., Burdette, E.C., Chirikjian, G., Fichtinger, G.: A robust fluoroscope tracking (ftrac) fiducial. Med Phys **32** (2005) 3185–98
7. Ma, Y., et al: An Invitation to 3-D Vision. Springer (2000)
8. Holt, R and et al: Uniqueness of solutions to three perspective views of four points. IEEE Trans. on PAMI **17** (1995) 303–307
9. Jain, A., Kyon, R., Zhou, Y., Fichtinger, G.: C-arm calibration - is it really necessary? In: MICCAI. (2005) LNCS 3749, 639–646
10. Kon, R., Jain, A., Fichtinger, G.: Hidden seed reconstruction from c-arm images in brachytherapy. In: IEEE ISBI. (2006) 526–29

A small size fluoroscope tracking (FTRAC) fiducial

Ameet K Jain, Iulian Iordachita, Gabor Fichtinger

Johns Hopkins University, USA

jain@cs.jhu.edu

Introduction: C-arm fluoroscopy is the most widely used intra-operative imaging modality in orthopedic surgery [1], but it presently lacks the ability for robust and easy quantitative guidance. Quantitative reconstruction needs to solve the problems of C-arm calibration (incl. image distortion) and C-arm pose tracking. Though the first is a well-studied problem, pose recovery on un-encoded C-arms is still a major technical problem.

Off-the-shelf tracking devices were a natural choice to achieve pose recovery. Unfortunately, auxiliary trackers are sometimes impractical for various reasons like expense, complexity and clinical feasibility (optical trackers require line of sight, while EM trackers suffer from field distortion). This led us to revisit fiducial-based tracking. The history of radio-opaque fiducials for C-arm tracking starts as early as 1987 (a review can be found in [2]), with the number of beads going as high as 613 for some of the preliminary works. Though the accuracies were acceptable, the fiducials tended to be too large and cumbersome to use. To make the fiducials feasible, recent publications have proposed smaller fiducials by compact bead placement, at the expense of a decrease in accuracy.

Our previous contribution was the development of a novel fluoroscope tracking (FTRAC) fiducial [2]. The FTRAC fiducial uses parametric curves, in particular, proposing the use of ellipses. 3D ellipses are constrained to always project as ellipses (closed parametric curves), they achieve sub-pixel image segmentation, and allow for a closed form mathematical framework. Thus the FTRAC fiducial not only offered a decrease in size, but also an enhancement in accuracy, allowing for precise object reconstruction. We have previously reported a 3x3x6 cm fiducial. In this paper, we propose to further decrease the size at a marginal loss in accuracy.

Methods: The primary contribution of this work is the use of small size parametric curves (ellipses & straight lines), in addition to spherical beads that have been used conventionally. The fiducial encodes six DOF in a single image by creating a unique view from any direction. Theoretically, three points or three straight lines or one ellipse with one point is sufficient for pose recovery. In practice, any redundancy will boost accuracy. We used an optimized constellation of six beads, three straight lines and two ellipses.

First the fiducial is segmented automatically. Now, given any pose of the C-arm, a 48 dimensional error metric is computed, accompanied by its corresponding Jacobian. This is done via obtaining an expected projection of the fiducial and comparing against its true segmentation. Next, a non-linear Gauss-Newton optimizer rapidly computes the pose of the C-arm. The algorithm converged without failure, in over tens of thousands of runs.

An un-optimized Matlab implementation takes around 3.5 s to run (2.4 GHz, Intel P4, 512MB RAM, Windows 2000; excl. image transfer and processing time). This indicates that an optimized C++ implementation would perform near real-time.

Results: An already reported larger version of the fiducial (3x3x6 cm) was previously tested using synthetic data and also experimentally on a precisely machined mechanical phantom. The simulation experiments suggested an accuracy of 0.2 mm in translation, while the phantom experiments exhibited an accuracy of 0.56 mm (0.33° in rotation and 0.7 mm for 3D reconstruction). Thus, simulations tended to underestimate the experimental error by about 60%.

We studied the sensitivity of pose recovery by scaling the diameter & length, but without significantly altering the relative configuration. To accommodate the bead/wire thickness, a minimum separation is maintained between neighboring features. Hence, the length of the fiducial is not scaled directly in proportion to the diameter. Though the length of a 30 mm diameter fiducial is 60 mm, the length of a 10 mm diameter fiducial is maintained at 35 mm (as compared to 20 mm).

The attached figure (0.5 mm simulated image segmentation error) plots the translation and rotation errors as a function of fiducial size. As expected, the translation error along the imaging direction was significantly higher than that parallel to the image plane. The rotation error on the other hand was similar for

all axes. The simulated translational tracking accuracy of a 10 mm diameter fiducial was 0.4 mm. Assuming the 60% underestimation from simulations, a 10 mm diameter FTRAC fiducial promises an accuracy of about 1 mm in clinical circumstances. Using the same underestimation, the 15 mm and 20 mm FTRAC fiducials are expected to offer 0.85 mm and 0.7 mm tracking accuracies, respectively. Phantom experiments to validate the same are currently underway.
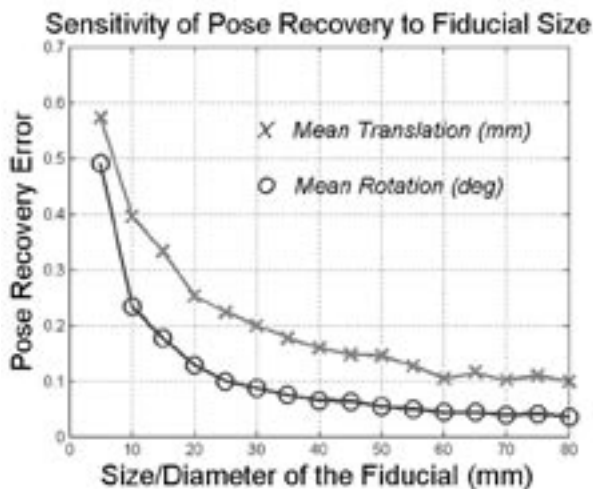
Dependence of C-arm tracking to fiducial size.

Discussion: A small fiducial for radiographic C-arm tracking was designed and computationally evaluated. The method promises accuracies similar to other tracking alternatives, and to be sufficiently robust for intra-operative quantitative C-arm fluoroscopy. Simulation studies indicate that the size can be reduced to 1x1x4 cm, while still maintaining an accuracy of about 1 mm.

There is a tradeoff between clinical feasibility (fiducial size) and desired accuracy, with a small sized FTRAC fiducial further facilitating a plurality of fluoroscopically guided procedures. The smaller sized FTRAC fiducial holds promise for tracking pencil-like devices such as needle guides, drill guides, etc (using a stationary C-arm). Furthermore, it has been shown recently that miscalibrated C-arms can still reconstruct 3D information with high accuracy [3] in several applications. Thus a short-range conventional EM or optical tool tracker, along with the FTRAC fiducial and an uncalibrated C-arm could still allow for accurate virtual fluoroscopy or 3D guidance with reconstruction.

References

[1] J. Yao, R. H. Taylor, R. P. Goldberg, R. Kumar, A. Bzostek, V. R. Van, P. Kazanzides, and A. Gueziec, "A c-arm fluoroscopy-guided progressive cut refinement strategy using a surgical robot," Comput. Aided Surg. 5, 373–390 (2000).

[2] A. Jain, T. Mustufa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (FTRAC) fiducial," Med. Phys. 32, 3185–3198 (2005).

[3] A. Jain, R. Kon, Y. Zhou, and G. Fichtinger, "C-arm calibration – is it really necessary?", MICCAI 2005, LNCS 3749, pp 639 – 646.

Experimental Setup: The accuracy of the tracking algorithm was validated on an X-ray fluoroscopy test bench. A modified anthropomorphic Rando™ phantom was placed in the field-of-view in order to emulate the background presented by a patient. Two sites on this phantom were used – the pelvis and the thorax. The pelvis presents a case of low signal while the thorax presents a cluttered background. The tool to be tracked was attached to a linear actuator and moved in known increments while capturing fluoroscopic images. The acquired image sequence was then sampled to generate new image sequences for testing the tracking algorithm. **Results:** The algorithm was robustly able to track the tool under a low signal in the abdomen, and under dense clutter presented by the ribs in the thorax. The rms error in each case was found to be as little as 0.7mm and 0.12mm, respectively. **Conclusion:** This study demonstrates the robustness of estimation under noise and clutter that can be achieved when tracking tools in a fluoroscopic sequence using a particle filtering approach.

WE-C-330A-02

Cone-Beam CT for Image-Guided Head and Neck Surgery: Assessment of Dose and Image Quality Using a C-Arm Prototype

M Daly*, [1] J Siewerdsen, [1, 2] D Moseley, [1] D Jaffray, [1, 2] (1) Ontario Cancer Institutue, Princess Margaret Hospital, Toronto, ON, Canada, (2) Department of Medical Biophysics, University of Toronto, Toronto, ON, Canada

Purpose: To quantify radiation dose to patient and staff during intraoperative cone-beam CT (CBCT)-guided head and neck surgery, and investigate 3D imaging performance as a function of dose. **Method and Materials:** Dose and image quality measurements were acquired on a mobile C-arm modified at our institution to provide flat-panel CBCT. Imaging dose was measured in a custom-built 16 cm cylindrical head phantom at four positions (isocenter, anterior, posterior, and lateral) as a function of kVp (80–120 kVp) and C-arm trajectory ("tube-under" and "tube-over" 180^o orbits). In-room exposure was measured at positions around the operating table and up to 2 m from isocenter. 3D image quality was assessed in CBCT reconstructions of an anthropomorphic head phantom containing contrast-detail spheres (11-103 HU; 1.6-12.7 mm). The contrast-to-noise ratio (CNR) was evaluated across a broad range of dose (0.6–23.3 mGy). **Results:** Dose in the 16 cm phantom (100 kVp; "tube-under" orbit) was 0.059 (isocenter), 0.022 (anterior), 0.10 (posterior), and 0.056 (lateral) mGy/mAs. Dose to the eyes (anterior) was reduced by a factor of 10 for "tube-under" versus "tube-over" orbits. In-room exposure for a typical CBCT scan (~10 mGy to isocenter) ranged from 33 mR at the anesthetist position, to <0.5 mR at 2 m from isocenter. CNR increased as the square root of dose, with excellent visualization of bony and soft-tissue structures in the anthropomorphic head phantom achieved at ~3mGy and ~10mGy, respectively. **Conclusion:** The prototype C-arm CBCT system demonstrates excellent visualization of bony and soft-tissue structures at dose levels low enough for repeat intraoperative imaging. High-performance image-guidance with respect to bony and soft-tissue anatomy was achieved at doses <~3mGy and <~10mGy, respectively. For guidance of head and neck surgery, significant dose sparing to the patient's eyes (a factor of 10) is achieved using a "tube-under" (rather than "tube-over") 180^o orbit.

WE-C-330A-03

Seed Segmentation in C-Arm Fluoroscopy for Brachytherapy Implant Reconstruction

S Vikal*[1], A Jain[2], A Deguet[2], D Song [3], G Fichtinger[2], (1) RWTH Aachen University of Technology, DE, (2) Johns Hopkins University, Baltimore, MD, (3) Johns Hopkins University School of Medicine, Baltimore, MD

Purpose: Intra-operative dosimetry in prostate brachytherapy critically depends on discerning the 3-D locations of implanted seeds. The accuracy of 3-D seed reconstruction step is, in turn, limited by the accuracy with which the position and orientation of individual implanted seed in the fluoroscopic images can be found. A method for robustly segmenting the seeds in fluoroscopic images is proposed here. **Methods and Materials:** The process of determining the locations and orientations of implanted seeds is sub-divided into three main steps. In the first step, the image is segmented by shape-size based morphological approach to eliminate background noise and do away with non-uniform brightness of the image, to get seed-like regions. These regions are either single seeds or overlapping multiple seed clusters. In the second step, the regions are

analyzed and classified definitively, in a two-phase statistical process coupled with information extraction from original intensity image, into two classes: single seed and overlapping multiple seed cluster. In the third step, the region belonging to overlapping multiple seed cluster is resolved into its constituent individual seeds through a simple and novel technique. **Results:** The proposed algorithm was tested on a set of ten clinical fluoroscopic images. The algorithm correctly determines the seeds with overall average of 99.57%. The clusters are not correctly resolved only in two images (2 clusters each, 1.7% and 1.6% of total seeds in respective implants). One false positive (noise labeled as seed) each is reported in two images, both the cases being where the tip of catheter appears to be of the size and shape of seed. **Conclusions:** The algorithm builds on an existing framework of morphological processing and provides further improvements in classification and cluster resolution. The algorithm appears to be robust and accurate despite the poor resolution of clinical images.

WE-C-330A-04

Effect of Projection Angles Used in Multi-View Reconstruction (MVR) Using Images From a Microangiographic (MA) Detector and An Image-Intensifier (II) System

V Patel*, A Kuhls, P Noël, A Walczak, C Ionita, R Chityala, R Tranquebar, H Rangwala, S Kasodekar, G Yadava, K Hoffmann, D Bednarek, S Rudin, Toshiba Stroke Research Center, SUNY at Buffalo, Buffalo, NY

Purpose: The sensitivity of a new 3D Multi-View Reconstruction (MVR) angiography technique to the projection angles used is evaluated by comparing 3D centerlines calculated from combinations of three projections acquired from two imaging systems with that from micro-Cone Beam CT (μCBCT), which is taken as truth. **Method and Materials:** A 3D centerline of a contrast-filled carotid vessel phantom was reconstructed from image data acquired using a custom-made μCBCT system with a microangiographic (MA) detector (45 μm pixels, 4.5 cm field-of-view (FOV)). Projection images of the same phantom were also acquired using the MA and an image intensifier (II) detector system (120 μm pixels, 4.5 in FOV) on a C-arm x-ray unit. The MVR technique was used to compute 3D centerlines for 12 combinations of projection angles. Each 3D MVR centerline was aligned with the μCBCT "true" 3D centerline using a Procrustes technique, and a root-mean-square (RMS) deviation was calculated. **Results:** The average RMS deviation for the MA-MVR centerlines is 25 μm with a standard deviation of 3 μm over the 12 different projection-angle combinations, whereas the average RMS deviation for the II-MVR centerlines is 41 μm with a standard deviation of 4 μm over these same combinations. The RMS deviation as a percent of the internal vessel diameter, 0.75 mm, is 3.3% for the MA and 5.5% for the II and appears to be independent of view selection. **Conclusion:** For the MVR technique, the improved resolution of the MA resulted in improved centerline determination compared to the II system. For both detectors, the selection of a particular projection set had little effect on the RMS centerline deviation. The low RMS deviations for both detectors indicate that the MVR technique can provide accurate 3D centerlines.

WE-C-330A-05

Segmentation of Radioactive Seed in 3D Ultrasound Images for Intraoperative LDR Prostate Brachytherapy

Z Wei*, London Health Sciences Center, and Robarts Research Institute, London, Ontario, Canada; A Fenster, Robarts Research Institute, and the University of Western Ontario, London, Ontario, Canada

Purpose: Develop and evaluate an algorithm to automatically localize implanted radioactive seeds in 3D ultrasound images for dynamic intraoperative low dose rate (LDR) brachytherapy procedures, in which all phases of the procedure are performed in one session to deal with variability in the current prostate brachytherapy. **Method and Materials:** Intraoperative seed segmentation in 3D TRUS images is achieved by performing a subtraction of the image before the needle has been inserted, and the image after the seeds have been implanted. The seeds are searched through a thresholding operation in a "local" space determined by the needle position and orientation information, which are obtained from a needle segmentation algorithm. To test this approach, 3D TRUS images of

Exhibit Hall F

Imaging
Moderated Poster Session
Moderated Poster - Area 3 (Imaging): Image Segmentation, Visualization, and Registration

TU-EE-A3-01

2D-3D Registration of Portal Images with the Planning CT for Detection of Patient Positioning Errors

G Mu*, P Xia, O Morin, UC San Francisco, San Francisco, CA

Purpose: To compare the use of 2D-3D automatic registration of portal images with the planning CT for detection of patient positioning errors and the use of 3D-3D registration of MVCBCT with the planning CT. **Method and Materials:** Two prototype programs were used to carry out 2D-3D and 3D-3D image registrations. To assess the accuracy and robustness of these programs, 25 sets of 2D portal images, 25 sets of megavoltage conebeam CT (MVCBCT) images with known positioning shifts were acquired. A planning CT of the RANDO was also acquired. The known shifts between these image sets were ranged from –17mm to 4 mm, -20mm to 5mm and –12mm to 6mm, with uncertainty of 4.278mm, 5.359mm, 3.396mm along the latitude, longitude and vertical directions **Results:** The average differences between 2D-3D method and the known shifts were -0.632±0.318 mm, -0.121±0.437 mm, -0.416±0.346 mm, compared to 3D-3D method of 1.487±0.342 mm, -0.127±0.528 mm, 0.083±0.48 mm along the latitude, longitude and vertical directions. The average differences between 2D-3D and 3D-3D image registration methods were 0.86±0.286 mm, -1.39±0.347 mm, -0.33±0.303 mm **Conclusion:** Both 3D-3D and 2D-3D registration methods can detect positioning errors within 1 mm. For a rigid body, 2D-3D method is sufficient.

Conflict of Interest: This project is partly funded by SIEMENS.

TU-EE-A3-02

Evaluation of a 2D-3D Registration Method for External Beam Radiation Therapy

H Jans*, A Syme, S Rathee, B Fallone, Cross Cancer Institute, Edmonton, Alberta

Purpose: To implement and validate a 2D-3D registration method for determining 3D patient position in external beam radiotherapy using orthogonal EPID images and megavoltage digitally reconstructed radiographs (MDRRs). To test the methods dependence on cost function, image pre-processing and parameter space sample density, and determine the dependence of registered rotations on setup translations and *vice versa*. **Method and Materials:** Orthogonal EPID image of a humanoid phantom in different poses (3D rotations and translations) were acquired in anterior-posterior and latero-lateral view. The EPID images were registered with a data base of orthogonal MDRRs, calculated as projection images through the phantom's CT data set at rotation angles within ±5°. Registration results were compared for three different cost functions (least-squares, cross-correlation and mutual information), different image pre-processing techniques (unsharp masking, histogram matching) and for isolated and combined rotations and translations. The influence of setup translations on registration results for rotations, and *vice versa*, was investigated and compared with a simple model. **Results:** Image pre-processing improves registration precision by more than a factor 2. Three dimensional translations were registered with better than 0.5 mm (one standard deviation) when no rotations were present. Three-dimensional rotations registered with a precision of better than 0.2° (1 SD) when no translations were present. Combined rotations and translations of up to 4° and 15 mm were registered with a precision of better than 0.4° and 0.7 mm respectively. Mutual information resulted in the most precise registration. Setup translations influence registered rotations, mostly following a simple theoretical model, but not *vice versa*. **Conclusion:** Precise registration requires image pre-processing and benefits from interpolation of the parameter space. Influence of object translation on registration of out-of-plane rotations can be significant; these "pseudo rotations" can be corrected using the theoretical model when only one projection image is used for registration (e.g. fluoroscopy).

TU-EE-A3-03

Automated Segmentation of Radiographic Fiducials for C-Arm Tracking

S Vikal*[1], A Jain[2], A Deguet[2], D Song[3], G Fichtinger[2], (1) RWTH Aachen University of Technology, DE, (2) Johns Hopkins University, Baltimore, MD, (3) Johns Hopkins University School of Medicine, Baltimore, MD

Purpose: Intraoperative quantitative C-arm fluoroscopy guidance depends on discerning the relative pose of images (pose recovery). A possible method is to use radiographic fiducials visible in fluoro images [1,2]. We propose a robust and fast method for segmenting fiducials designed for brachytherapy applications. **Methods and materials:** The fiducial contains points, lines and ellipses made from BBs and wires[1]. The algorithm integrates the a-priori knowledge of fiducial's mechanical construction in a cleverly devised workflow. The BB segmentation is achieved using morphological top-hat transform. This information serves as a heuristic input to line segmentation realized by a curve tracing algorithm which operates on edge image, followed by augmenting information from intensity image. Once the lines are segmented, this information feeds to the ellipse extraction step. For ellipse segmentation, intensity image is morphologically processed to eliminate background noise, followed by elimination of BB-s and lines from the information obtained in prior steps. The resulting image consists of only ellipse segments. A fast variation of Hough transform is used to rectify the full ellipse from the segments. **Results:** The fiducial algorithm identified all the features (BBs, lines and ellipses) visible to human eye in all ten clinical images. Next the accuracy of fiducial segmentation was assessed numerically by feeding the results to the pose recovery algorithm of [1]. The fiducial was moved on an accurate mechanical platform (as ground truth) while the C-arm was stationary. We reconstructed the relative poses with an accuracy of 1.2 mm in translation and 0.3 degrees in rotational based on the segmented fiducials. **Conclusions:** The algorithm makes effective use of a-priori knowledge and combines the techniques of morphological segmentation, curve tracing, and Hough transform, resulting in a novel curve segmentation strategy.

References:

[1] Jain et al, Med Phys 32(10):3185-98
[2] Zhang et al, Phys Med Biol 49: 335–345

TU-EE-A3-04

Massive Training Artificial Neural Network (MTANN) to Reduce False Positives Due to Rectal Tubes in Computer-Aided Polyp Detection

K Suzuki*[1], H Yoshida[2], J Nappi[2], S Armato[1], A Dachman[1], (1) University of Chicago, Chicago, IL, (2) Massachusetts General Hospital, Boston, MA

Purpose: One limitation of current computer-aided detection (CAD) of polyps in CT colonography is a relatively large number of false positives. Rectal tubes are a common source of false positives and may distract the reader from less common polyps in the rectum. Our purpose was to develop a three-dimensional massive-training artificial neural network (3D MTANN) for reduction of false positives due to rectal tubes generated by a CAD scheme. **Material and Methods:** Our database consisted of CT colonography of 73 patients, scanned in both supine and prone positions. Fifteen patients had 28 polyps (15 polyps: 5-9 mm; 13 polyps: 10-25 mm). These cases were subjected to our previously reported CAD scheme that included shape-based detection of polyps and reduction of false positives with a Bayesian neural network. With this scheme, 96.4% (27/28) by-polyp sensitivity with 3.1 (224/73) false positives per patient was achieved. To eliminate false-positive rectal tubes, we developed a 3D MTANN that was trained to enhance polyps and suppress rectal tubes. **Results:** In the output volumes of the trained 3D MTANN, various polyps were represented by distributions of bright voxels, whereas rectal tubes appeared as darker voxels. The 3D MTANN removed all 20 false-positive rectal tubes produced by our original CAD scheme without removing any true positives. To evaluate the overall performance, we applied the 3D MTANN to the entire database containing 27 polyps (true positives) and 224 non-polyps (false positives). The 3D MTANN eliminated 33% (73/224) of non-polyps without removal of any true positives in an independent test.**Conclusion:** The 3D MTANN was able to improve the false-positive rate of our original CAD scheme from 3.1 to 2.1 false positives per patient, while an original by-polyp sensitivity of 96.4% was maintained. **Conflict of Interest:** HY, SGA: shareholders, R2 Technology, Inc.

HIDDEN SEED RECONSTRUCTION FROM C-ARM IMAGES IN BRACHYTHERAPY

Ryan C. Kon, Ameet Kumar Jain, and Gabor Fichtinger

Johns Hopkins University
Engineering Research Center for Computer Integrated Surgical Systems and Technology
{jain, gabor}@cs.jhu.edu

ABSTRACT

There has been a pressing clinical need for adaptive intra-operative dosimetry in the delivery of prostate brachytherapy implants. The missing prerequisite is the robust matching of the seeds across multiple C-arm images. This is further aggravated since seeds are invariably hidden in each image. We present a solution to recover these hidden seeds in this paper. A network flow formulation of the problem is proposed, where the desired solution is obtained (in polynomial time) by computing the flow with minimum cost. Phantom experiments show that using four X-ray images, on an average 99.8% of the seeds are recovered correctly, while simulations indicate that our algorithm is robust to segmentation errors of up to $1\ mm$ and hidden seed rate of at least 8%. The results show strong feasibility and clinical data collection is currently underway.

1. MOTIVATION AND BACKGROUND

With an approximate annual incidence of 220,000 new cases and 33,000 deaths, prostate cancer continues to be the most common cancer in men in the United States. The definitive treatment modality for low risk prostate cancer is permanent brachytherapy, which is performed on approximately 40,000 patients each year. In this treatment a large number of small ($\sim 1x5\ mm$) radioactive capsules are implanted into the prostate to kill the cancer by emitting radiation. According to a comprehensive review by the American Brachytherapy Society [1], *the preplanned technique used for permanent prostate brachytherapy has limitations that may be overcome by intraoperative planning*. However, the report continues, *the major current limitation of intraoperative planning is the inability to localize the seeds in relation to the prostate*.

Brachytherapy is typically performed transperineally under real-time transrectal ultrasound (TRUS) guidance and C-arm fluoroscopy is often used for gross visual observation of the implant. TRUS usually provides adequate imaging of soft tissue anatomy, but it fails to visualize implanted seeds, while C-arm fluoroscopy can visualize seeds, but not soft tissues. By reconstructing the implanted seeds from C-arm fluoroscopy and registering them to ultrasound, intra-operative dosimetry becomes possible.

3D coordinates of the implanted seeds can be calculated by resolving the correspondence of seeds across multiple X-ray images. Five major obstacles need to be overcome: (a) C-arm calibration; (b) C-arm pose tracking; (c) Seed segmentation; (d) Seed matching and reconstruction; and (e) Registration between C-arm and TRUS. While adequate techniques are available for most of these problems [2] and several groups have published results supporting C-arm fluoroscopy for intra-operative dosimetry [3], this technique has yet to become a standard of care in hospitals. The last technological barrier appears to be robust matching and reconstruction of seeds across multiple C-arm images. This problem is still unsolved because implants are usually dense and all the seeds are hardly ever visible in any one image. Even in the case of hand segmentation, perfectly overlapping seeds cannot be segmented ($\sim 2-6\%$). This phenomenon is often referred to as the issue of "hidden seeds" and is the focus of this paper.
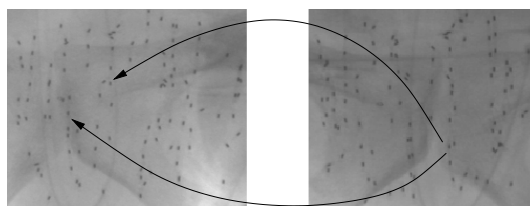
Fig. 1. Example of a hidden seed not detectable by the human eye (right) matching to two distinct seeds (left).

Many contemporary works have made a simplifying assumption that all the seeds are visible, which makes these algorithms infeasible for clinical use. Among previous works, Fast-CARS was extended to incorporate hidden seeds, but the new algorithm reconstructed a greater number of seeds than were actually present [4]. Another variant was proposed [5] by ordering the seeds using the epipolar constraints. Unfortunately, the algorithm required co-planar images (co-linear X-ray sources) and could not reconstruct undetected seeds if they existed in the same search restriction band and did not

extend to multiple images. An intensity-based method using tomosynthesis [6] and using Hough trajectories [7] has also been proposed. However they require an unfeasibly large number of images to achieve a stable reconstruction, yet do not offer accuracies better than $1mm$. Another technique [8] optimized on seed positions and camera parameters, by generating simulated images and iterating them until they match the observed images. This optimization method is prone to fall into local minima and was tested only on clean simulated images. Significant works as they are, the problem merits further research to produce a clinically viable solution.

We have previously proposed matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL) [9]. In this paper, we extend MARSHAL to also tackle the hidden seed problem.

2. METHODS AND MATERIALS

We assume that the seeds are 3D points and convert the seed-matching problem to a network-flow-based combinatorial global optimization. In this formulation, any correspondence of the seeds is represented by an appropriate flow through the network. The goal is to find the flow with the minimum cost. In the case with no hidden seeds, two images do not have a unique solution, while using three or more images makes the problem NP-hard [9]. Since the case with hidden seeds is a generalization of the earlier problem, it is also NP-hard. This proves that a globally optimal polynomial time algorithm is not possible.

In contrast to the previously proposed heuristic methodologies, MARSHAL solves the problem in a more mathematically rigorous framework of combinatorial optimization. This formal approach allows better control of the behavior of the algorithm, as well as consideration of seed groups as a whole (global optimization) instead of analyzing seed groups with heuristic rules (local optimization).

A Network-flow-based Formulation: Let N be the number of seeds inserted and N_1, N_2, and N_3 be the number of seeds actually segmented in the acquired C-arm images I_1, I_2, and I_3. Let s_{ij} be the position of the i^{th} seed in the j^{th} image. We construct a directed network as shown in Fig 2.

Sets A, B, C, and D, each containing N_1, N_2, N_3, and N_1, respectively, represent images I_1, I_2, I_3, and I_1, respectively. There are no edges between nodes of the same set, but there are directed edges (links) connecting every node in set A (left) to set B, set B to set C, and set C to set A (right). There are also links from the source S to every node in set A (left), and similarly from set A (right) to the sink T. A flow of value N originates at S and ends at T. Because any of the sets can contain fewer than N nodes, each link from S to A or from A to T must allow for a flow greater than 1. A flow of 2 at a node implies that this seed is a hidden seed. To enforce the constraint that each seed in image 2 & 3 is chosen at least

once but not more than twice, we add dummy internal links. Thus, we allow a maximum/minimum flow of 2/1 through each of the internal links and through nodes connected to the source and sink. The links connecting the sets allow for a maximum/minimum flow of 1/0.
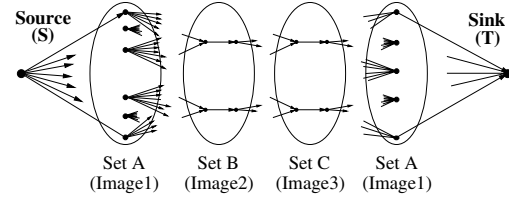

Fig. 2. The flow network used to solve the matching problem.

The problem now reduces to efficiently computing a flow of N. To determine the optimal solution, we need to assign a cost $\mathcal{C}_{ij}$ to the link connecting seed s_{i1} to seed s_{j2}. $\mathcal{C}_{ij}$ represents the likelihood of seed s_{i1} matching seed s_{j2}, with the cost being 0 if they match perfectly and ∞ (infinity) if they do not match at all, with extensive discussions available [9]. Any feasible flow in the network has a net cost of $\sum \mathcal{C}_i f_i$, where f_i is the flow in link i and $\mathcal{C}_i$ is the cost of sending a unit flow along that link. Thus the seed-matching problem is reduced to finding the flow with minimum cost.

The minimum-cost flow can be computed using many polynomial and pseudo-polynomial algorithms [10]. We implemented the cycle cancelling algorithm to compute min-cost flow, which starts with a feasible flow and then searches for negative cost cycles in the residual network and pushing a flow along the negative cycle. This is iterated until there is no negative cost cycle detected, at which point the minimum cost flow has been computed. Negative cost cycles can easily be computed using the Bellman-Ford algorithm, which runs in $O(Edges * Vertices)$ time. Thus the run-time for the min-cost flow evaluation is $O(cost\ of\ the\ initial\ flow) \times O(Edges*Vertices)$, which will be pseudo-polynomial. Although there are strongly polynomial time algorithms available, the cycle-cancelling algorithm proves to be sufficiently fast.

Overview of the Algorithm: Extended MARSHAL is illustrated in Fig 3. Since the problem is NP-hard the network is not able to constrain the hidden seeds in image 1. Nevertheless, hidden seeds in image 2 & 3 are very well constrained. Thus the min-cost flow nearly solves the problem, except for a few cases which exhibit some ambiguity. These cases usually incarnate themselves as self-consistent complete subsets of 2-5 seeds. Each of these subsets can be independently solved using an extremely fast brute-force type of algorithm.

In the case that the size of any subset is large, the subset can be resolved by constructing a miniature network and solving for the min-cost flow. It is unusual to encounter small subsets which have hidden seeds in all three images, and hence
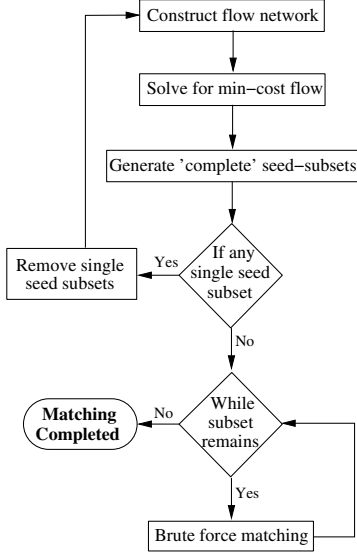
Fig. 3. The flowchart for extended MARSHAL.

Fig. 4. (a) Sensitivity to segmentation error. (b) Sensitivity to hidden seed percentage.
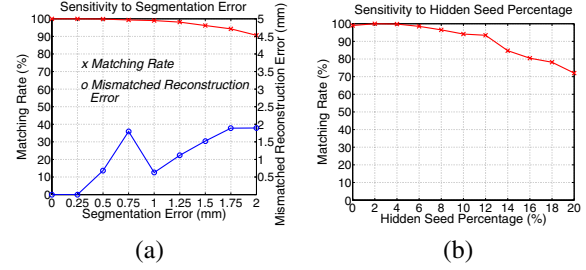
these subsets can be resolved correctly by switching the order of the images. In the rare case that all three images in the subset have hidden seeds, then a choice of set A can be made arbitrarily and the network solved. Seeds solved as hidden in sets B/C can then be replicated and used as the outer set and the network solved again.

The robustness of the reconstruction can be significantly improved by taking another X-ray image, and thus we extend the algorithm to work with any number of images. The problem is well conditioned due to the epipolar constraints, so when four images are available, three images can be chosen as the main images, while one can be chosen as an assisting image (one with the most number of hidden seeds) affecting only the costs in the main network. This approach can be extended to any number of images. Thus extended MARSHAL always recovers the hidden seeds and resolves the correspondences in polynomial time (typically close to $O(N^3)$ [9]).

3. RESULTS AND DISCUSSION

Simulations: Data was generated to simulate a $55\ cc$ prostate with a seed density of $2.0\ seeds/cc$. The algorithm was run on three different datasets using combinations of four images, varying segmentation error from $0-2\ mm$ in increments of $0.25\ mm$. To test the sensitivity to the number of hidden seeds, the algorithm was also run varying the hidden seed percentage from $0-20\%$ in increments of 2%. Averaged results are shown in Fig 4.

Phantom experiments: Experiments were conducted on a precisely fabricated acetol seed phantom. The FTRAC fiducial [2] was used to track the C-arm, and was attached to the phantom (Fig 5). The phantom comprises of twelve $5\ mm$

thick slabs, each having one hundred holes with $5\ mm$ spacing. Known implant constellations were created, with the number of seeds ranging from 40 to 100 in increments of 15, while keeping the seed density at $1.56\ seeds/cc$.

For a given constellation, 6 images within a 20^o cone around the AP-axis were taken using an *uncalibrated Philips Integris V3000* and dewarped using the pin-cushion test. Accurate ground truth for matching was computed from the known 3D seed locations. Matching was achieved with three and four images. All seeds closer than $1.2\ mm$ were called hidden. Averaged results are displayed in Table 1.
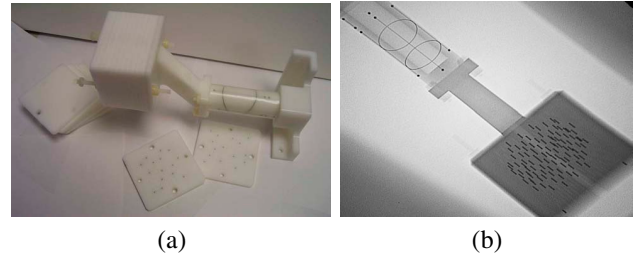
Fig. 5. (a) An image of the phantom attached to the FTRAC fiducial. (b) A typical X-ray image of the combination.

Discussion: The simulation results show that the algorithm can nearly perfectly match all the seeds even with segmentation errors up to $1\ mm$. While increasing the segmentation error further decreases the matching percentage, at $2\ mm$ error, the matching rate is still over 90%, with the reconstruction error of the mismatched seeds remaining below $2\ mm$. When the percentage of hidden seeds is varied, the algorithm can robustly match when the hidden seed percentage is at least 8%. Because the datasets for this case were generated by creating a threshold for closeness based only on the hidden percentage, significant error was introduced as the percentage increased, since the threshold became unrealistically large, creating much lower matching percentages than would normally be seen.

For the phantom data, using three images gives a good matching rate, but mismatched seeds reconstruct with a high error. Using four images gives substantially better results,

| | | Number of Seeds | | | | | | | | | | |
| | | 3 Images | | | | | | 4 Images | | | | |
		40	55	70	85	100		40	55	70	85	100
Matching Rate (%)	Match	98.3	100	99.9	98.2	98.0		100	100	100	99.9	99.5
Reconstruction Error (mean)	Match	0.64	0.50	0.67	0.71	0.77		0.64	0.39	0.63	0.76	0.82
Reconstruction Error (STD)	Match	0.27	0.27	0.31	0.26	0.34		0.25	0.21	0.29	0.27	0.28
Reconstruction Error (mean)	Mismatch	7.46	-	0.80	12.47	4.99		-	-	-	14.13	1.73
Reconstruction Error (worst)	Mismatch	9.02	-	0.80	17.34	8.38		-	-	-	14.13	1.89
Reconstruction Error (relative)	All	0.52	0.32	0.40	0.54	0.53		0.32	0.30	0.38	0.34	0.41

Table 1. Performance on phantom data.

with nearly perfect matching and a mostly low reconstruction error for mismatched seeds. Since it is readily feasible to obtain a fourth image in a clinical setting, our implementation is extremely viable for intra-operative dosimetry.

Conclusions and Future Work: In contrast to other proposed methods, we have formalized the seed matching problem and have extended a previously proposed polynomial time algorithm (MARSHAL) to resolve hidden seeds. A MATLAB 7 implementation runs in under 20 s in a typical implant using any number of images. Using 4 images, it matched over 99.8% of the seeds. Simulations indicate that MARSHAL is robust to various parameters. It can reconstruct an implant when three or more images are used, with a robustness, precision, and speed that promises to be sufficient to support intraoperative dosimetry in prostate brachytherapy.

MARSHAL is being combinatorially improved to increase the matching rate when only three images are used and also to reject spuriously segmented seeds. Moreover, other sources of information like seed orientation are being added to the cost metric to increase robustness. Clinical data is also being collected for further analysis.

4. REFERENCES

[1] S Nag et al, "Intraoperative planning and evaluation of permanent prostate brachytherapy: report of the american brachytherapy society," *Int J Radiat Oncol Biol Phys*, vol. 51, no. 5, pp. 1422–30, Dec 2001.

[2] A Jain et al, "A robust fluoroscope tracking (ftrac) fiducial," *Med. Phys.*, vol. 32, no. 10, pp. 3185–98, Oct 2005.

[3] D Todor et al, "Intraoperative dynamic dosimetry for prostate implants," *Phys Med Biol*, vol. 48(9), pp. 1153–71, May 7 2003.

[4] Y Su et al, "Prostate brachytherapy seed localization by analysis of multiple projections: identifying and addressing the seed overlap problem," *Med Phys*, vol. 31(5), pp. 1277–87, May 2004.

[5] S Narayanan et al, "Three-dimensional seed reconstruction from an incomplete data set for prostate brachytherapy," *Phys Med Biol*, vol. 49(15), pp. 3483–94, Aug 2004.

[6] I Tutar et al, "Tomosynthesis-based localization of radioactive seeds in prostate brachytherapy," *Med Phys*, vol. 30(12), pp. 3135–42, Dec 2003.

[7] S Lam et al, "Three-dimensional seed reconstruction for prostate brachytherapy using hough trajectories," *Phys Med Biol*, vol. 49(4), pp. 557–69, Feb 2004.

[8] M Murphy et al, "Demonstration of a forward iterative method to reconstruct brachytherapy seed configurations from x-ray projections," *Phys Med Biol*, vol. 50, pp. 2715–37, Jun 2005.

[9] A Jain et al, "Matching and reconstruction of brachytherapy seeds using the hungarian algorithm (marshal)," *Med Phys*, vol. 32, no. 11, pp. 3475–92, Nov 2005.

[10] R Ahuja et al, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall; 1 edition, 1993.

Incorporating seed orientation in brachytherapy implant reconstruction

Yu Zhou*[a], Ameet K. Jain[b], Gregory S. Chirikjian[c], Gabor Fichtinger[b,c]
[a]Department of Mechanical Engineering, State University of New York at Stony Brook, Stony Brook, NY 11794;
[b]Department of Computer Science, The Johns Hopkins University, Baltimore, MD 21218;
[c]Department of Mechanical Engineering, The Johns Hopkins University, Baltimore, MD 21218.

ABSTRACT

Intra-operative quality assurance and dosimetry optimization in prostate brachytherapy critically depends on the ability of discerning the locations of implanted seeds. Various methods exist for seed matching and reconstruction from multiple segmented C-arm images. Unfortunately, using three or more images makes the problem NP-hard, i.e. no polynomial-time algorithm can provably compute the complete matching. Typically, a statistical analysis of performance is considered sufficient. Hence it is of utmost importance to exploit all the available information in order to minimize the matching and reconstruction errors. Current algorithms use only the information about seed centers, disregarding the information about the orientations and length of seeds. While the latter has little dosimetric impact, it can positively contribute to improving seed matching rate and 3D implant reconstruction accuracy. It can also become critical information when hidden and spuriously segmented seeds need to be matched, where reliable and generic methods are not yet available. Expecting orientation information to be useful in reconstructing large and dense implants, we have developed a method which incorporates seed orientation information into our previously proposed reconstruction algorithm (MARSHAL). Simulation study shows that under normal segmentation errors, when considering seed orientations, implants of 80 to 140 seeds with the density of 2.0- 3.0 seeds/cc give an average matching rate >97% using three-image matching. It is higher than the matching rate of about 96% when considering only seed positions. This means that the information of seed orientations appears to be a valuable additive to fluoroscopy-based brachytherapy implant reconstruction.

Keywords: Prostate brachytherapy, implant reconstruction, seed matching, C-arm fluoroscopy

1. INTRODUCTION

With an approximate annual incidence of 220,000 new cases and 33,000 deaths, prostate cancer continues to be the most common cancer in men in the United States[1]. For several decades, the definitive treatment for low risk prostate cancer was radical prostatectomy or external beam radiation therapy[2], but low dose rate permanent seed brachytherapy (shortly brachytherapy) today can achieve virtually equivalent outcomes[3,4]. The success of brachytherapy (i.e., maximizing its curative force while minimizing its co-morbidity) mainly depends on our ability to tailor the therapeutic dose to the patient's individual anatomy. In contemporary practice, however, implant planning is based on idealistic preplanned seed patterns. 15 years of clinical practice has clearly demonstrated that it is not achievable in the actual human body. According to a comprehensive review by the American Brachytherapy Society[5], the preplanned technique used for permanent prostate brachytherapy has limitations that may be overcome by intraoperative planning. At the same time, the major current limitation of intraoperative planning is the inability to localize the seeds relative to the prostate. There are excellent algorithmic and computational tools available today to optimize a brachytherapy treatment plan intraoperatively, thereby allowing for improved dose coverage. These methods, however, critically require that the exact 3D locations of the implanted seeds are precisely known with respect to the patient's anatomy.

Prostate brachytherapy is almost exclusively performed under transrectal ultrasound imaging (TRUS) guidance. While TRUS provides adequate imaging of the soft tissue anatomy, it does not allow for robust localization of the implanted brachytherapy seeds. Many researchers have tried to segment the seeds from TRUS images by linking seeds with spacers[6], using X-rays to initialize segmentation[7], using vibro-accoustography[8] or transurethral ultrasound[9] as a new imaging modality, or segmenting them directly[10]. But even when meticulously hand-segmented, up to 25% of the seeds

may remain hidden in ultrasound[11]. This necessitates the use of some other imaging method in intraoperative seed localization.

The application of C-arm fluoroscopy in brachytherapy originates when it was first used as a solo guidance modality[12]. Shortly after TRUS emerged as a primary image guidance modality, fluoroscopy became a secondary tool for gross visual observation. Mobile C-arms are ubiquitous in contemporary prostate brachytherapy, with approximately 60% of the practitioners using it for qualitative implant analysis in the operating room[13]. It is considered as the gold standard for intraoperative visualization of brachytherapy seeds. While several groups have published protocols and clinical outcomes favorably supporting C-arm fluoroscopy for intraoperative dosimetric analysis[10,14–22], this technique is yet to become a standard of care across hospitals.

The ability to reconstruct and register the implanted seeds, which are visible in fluoroscopy, to soft tissue anatomy, which is visible in TRUS, intraoperatively would allow us to make immediate provisions for dosimetric deviations from the optimal implant plan. At the same time, quantitative use of fluoroscopy for dosimetric analysis has been hampered by a series of unresolved technical problems. The five major obstacles toward intraoperative dosimetry are: (a) C-arm distortion correction and calibration, (b) C-arm pose tracking, (c) seed segmentation, (d) seed matching and reconstruction, and (e) registration of C-arm to TRUS images.

Significant efforts have been made toward computational fluoroscopy guidance in general surgery, developing various tools for distortion correction and calibration[23,24]. However, C-arms available in most hospitals do not have encoded rotational joints, so one never knows where the fluoro shots are coming from relative to one another. We have addressed this issue by designing a fluoroscope tracking (hence-forth FTRAC) fiducial, which is a radiographic fiducial system creating a unique projection image from each direction[25]. Various methods partially dealing with C-arm calibration in brachytherapy have also been proposed[26–28], while some others have suggested that it is redundant[29].

Attempts have also been made to relate fluoroscopic images to soft tissue anatomy[10,17,30–34]. Nevertheless, further research is merited since existing methods are susceptible to various kinds of errors. We addressed this issue by the use of the FTRAC fiducial[25]. It is capable of not only tracking the C-arm, but also registering the C-arm to TRUS by a predetermined placement.

Methods are available for automatic seed segmentation[15,35–38]. 3D coordinates of the implanted seeds can now be calculated from multiple X-ray images upon resolving the correspondence of seeds. Formalization of the seed-matching problem results in a high complexity search space of the order 10^{150} and 10^{300}, from two and three fluoroscopic images, respectively. Hence previously proposed seed-matching approaches have predominantly been heuristic explorations of the search space, with no theoretical assurance on the accuracy of the answer. The early attempts toward seed matching used three coplanar images (coplanar images are those where the implant and the three X-ray sources are approximately in the same plane)[39-41]. The images were divided into variable width bands, formed by comparing coordinates along the rotation axis. Furthermore, in order to make the bands, it was assumed that the seeds are near the iso-center of the C-arm or at least have similar magnifications in all the images. These methods are prone to calibration errors and become ineffective as the number of seeds increases. These ideas were further extended by accommodating for patients' motion[42], and yet all the seeds could not be reliably reconstructed. Further geometrical constraints were imposed by assuming that some of the seeds are in a straight line[43] or on quadratic curves[44], which is generally not true due to seed migration.

The first step toward mathematical formalization came with the construction of a cost matrix[45], where exhaustive matching gave the lowest cost solution. Though it eliminated extraneous assumptions, it required impractical computational resources. A greedy randomized algorithm, tested with various cost metrics, was suggested to reduce the runtime[46]. This method gives a different output for each run and is typically iterated a few hundred times, choosing the sequentially lowest cost. Though this method might provide an answer close to the correct match, its randomization does not make any claim on the number of iterations required for the final answer. Fast-CARS is another variant[47], which significantly improved the computational complexity, where for each cost matrix an exhaustive match can be performed to obtain the best possible matches. It reduced the run-time from $O((N!)^2)$ to $O((A!)^2)$, where A is the average number of seeds in the band. Though it made the search faster, it still ran in exponential time. For example, if A=10 then the number of computations would still be as high as $O(10^{14})$.

Independently, a set of heuristic rules[16], which attempted to reduce misclassifications, were suggested for seed matching. Simulated annealing[36] was proposed as an alternate technique to reach the global minimum. Another technique[28] was proposed, which optimizes seed positions and camera parameters by generating simulated images and iterating them until they match the observed images. These optimization methods are prone to local minima, and may not be able to recover from them. A statistical simulation of seed reconstruction uncertainty was conducted[48], but did not address the problem of seed matching. CT and MRI based techniques[49,50] were also proposed, but cannot be used intraoperatively, and have poor resolution in the axial direction.

The matching problem is also prevalent in the computer vision community, where 2D points are tracked and reconstructed to compute motion. Researchers have tried to use noniterative greedy algorithms[51], also incorporate spurious and hidden points[52,53]. Occlusion itself has also been a known problem[54]. These algorithms were optimized for a dense set of moving points, while specialized algorithms were used for sparse matchings[55,56], which can also be used in pattern recognition across images[57]. These algorithms are usually catered to achieve real-time performance, as compared to a complete matching, and hence do not appear to be appropriate in a medical application.

In our recent work[59], a new theoretical framework for seed matching has been introduced. The framework tackles issues of optimality, and presents a practical algorithm. Moreover, the framework guarantees a polynomial runtime of $O(N^3)$ on the algorithm, an improvement over previous methods. Seed matching and reconstruction is done by using the Hungarian algorithm (MARSHAL). Simulations and phantom experiments show that MARSHAL is not sensitive to image separation, seed density, the number of seeds, and C-arm calibration, and also robust enough to segmentation, C-arm pose, and distortion. It can reconstruct an implant when three or more images are used, with enough robustness, precision, and speed to support intraoperative dosimetry in prostate brachytherapy. It can also be used as a general purpose correspondence algorithm in many synergistic problems.

In seed matching and reconstruction, it is of utmost importance to exploit all the available information in order to minimize the matching and reconstruction errors. However, current algorithms, including MARSHAL, use only the information about seed centers, disregarding the information about the orientations and length of seeds. While the latter has little dosimetric impact, it can positively contribute to improving seed matching rate and 3D implant reconstruction accuracy. It can also become critical information when hidden and spuriously segmented seeds need to be matched, where reliable and generic methods are not yet available. Expecting orientation information to be useful in reconstructing large and dense implants, in this paper we propose a method which incorporates seed orientation information into MARSHAL.

This paper is organized in the following way. In Sec. 2, we will describe the proposed method which combines the information of seed orientations with that of seed positions. In Sec. 3, we will discuss the simulation study and the results. In Sec. 4, we will conclude our current work and discuss the future work.

2. METHODOLOGY

In this paper we further improve the matching rate of MARSHAL by including the information of seed orientations into the cost matrix. We convert the seed-matching problem to network-flow-based combinatorial optimization. In this formulation, any correspondence of the seeds is represented by an appropriate flow through the network. This formal approach allows better control of the behavior of the algorithm, considering the set of seeds in global optimization instead of local optimization as heuristic rules do.

2.1 A generic network-flow-based formulation
A network flow formulation is created for seed matching problem, where any flow in the network would represent a matching, and the desired solution is the flow with minimum cost. Let N seeds be implanted, and C-arm images I_1, I_2 be acquired. Let s_{ij} be the position of the i^{th} seed in j^{th} image. We construct a directed network as shown in Fig. 1. Sets A and B, each with N nodes, represent the two images I_1 and I_2. While there are no edges within each set, directed edges run from all vertices in set A to all vertices in set B. There are N edges at source S, each edge connecting to a node in A. Similarly each node in B is connected to sink T. The flow originates at S and ends at T, with each edge allowing a flow of value 1 or 0, where 1 means that the edge is selected and 0 means that it is not. The problem is to find a flow in the

network that can achieve a total flow of value N. To have a net flow of N, each edge connecting to either the source or the sink has to support a flow 1. Now by the conservation of flow at each node, every node in set A will have to dispatch a unit flow to some node in set B. Moreover, each node in set B can accept only a unit flow, because any extra flow cannot be passed on to T, and any deficiency would mean that T does not have a total flow of N units. The set of all edges with nonzero flow provide a feasible matching.
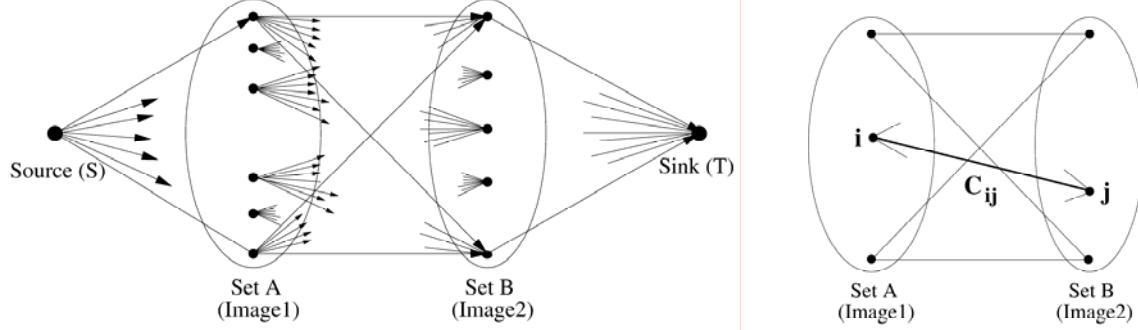

Figure 1: Network flow formulation for the seed matching problem   Figure 2: Two-image seed matching as the assignment problem

Under no constraint, the network flow problem has N! solutions, corresponding to N! feasible flows. To obtain the optimal solution, the edge connecting seed s_{i1} to seed s_{j2} is assigned a cost C_{ij}. The cost C_{ij} represents the likelihood of seed s_{i1} matching seed s_{j2}, with the cost being 0 if they match perfectly and ∞ (infinity) if they do not match at all. Any feasible flow has a net cost associated with it, the value of which is $\sum_{i=1}^{n}\sum_{j=1}^{n} C_{ij} f_{ij}$, where f_{ij} is the flow in edge ij, and C_{ij} is the cost of sending a unit flow along that link. Thus the seed matching problem is reduced to finding the flow with minimum cost, and can be written as

$$\min \sum_{i=1}^{N}\sum_{j=1}^{N} C_{ij} f_{ij}, \quad where \quad f_{ij} \in \{0,1\}, \quad \sum_{i=1}^{N} f_{ij} = 1 \, \forall i \quad and \quad \sum_{j=1}^{N} f_{ij} = 1 \, \forall j. \tag{1}$$

The network flow in Fig. 1 can easily be extended to multi-image cases by adding more sets of nodes and directed edges between S and T.

2.2 Seed matching for two and three images
When all the seed locations in the two images are known, the minimum-cost maximum-flow formulation reduces further to the specific problem of minimum-weight matching in bipartite graphs, also known as the assignment problem (Fig. 2). The problem is to find a minimum weight subset of edges such that all the vertices are covered exactly once. The assignment problem is solved in $O(N^3)$ runtime by using the Hungarian algorithm[60–62]. Since the algorithm is well-known, we just give an outline here. The N×N cost matrix C is constructed. The objective is to choose exactly one element from each row and column such that the sum of the elements has the lowest value. Thus, an equivalent matrix having at least one zero in each row and column is obtained by subsequent subtractions using the smallest element in each row and column. This matrix is used to find a selection of zeros such that each row and column has exactly one zero. If it exists, then it provides the minimum-cost matching. If it does not exist, a line covering procedure is used to adjust the matrix and generate zeros in useful locations. The locations of all zeros provide the minimum-weight matching. Thus the Hungarian algorithm provides the matching with the lowest possible cost.

Our previous work has proved that a robust seed matching and reconstruction requires at least three images[59]. In the three-image case, the network flow problem becomes a tripartite matching problem. We proposed a practical solution for matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)[59]. MARSHAL projects the original tripartite problem into three distinct bipartite problems by the appropriate projection of the costs. Each bipartite problem can be solved in $O(N^3)$ runtime. The solutions of the bipartite matchings are then combined to obtain a solution to the original tripartite problem. A detailed description of MARSHAL algorithm and its theoretical foundation can be found in Ref. 59. A discussion of applying MARSHAL to seed matching from four or more images can also be found there. To avoid redundancy, here we only list the flowchart for MARSHAL (Fig. 3).

2.3 Seed reconstruction

In principle, given three images, the 3D location of a seed can be computed by first defining the equation of the three straight lines that join each projection with its respective X-ray source, and then calculating the intersection point of the three straight lines. However, due to various errors, these straight lines never intersect, forcing us to compute a symbolic 3D intersection point. The symbolic intersection is typically defined as the global minimum of an error function. We have proposed a simple and quick method that minimizes the L_2 norm of Euclidian distance from the intersection point to the lines[59]. Here we assume that for a point P in space, there are m corresponding straight lines (Fig. 4). Line i (l_i) joins the projection of P in image i, p_i, with its X-ray source. The unit directional vector of line i is (a_i, b_i, c_i). Then the 3D location of the symbolic intersection point can be calculated as

$$P = (\sum_{i=1}^{m} A_i)^{-1} \sum_{i=1}^{m} A_i p_i , \qquad (2)$$

where

$$A_i = \begin{bmatrix} b_i^2 + c_i^2 & -a_i b_i & -a_i c_i \\ -a_i b_i & a_i^2 + c_i^2 & -b_i c_i \\ -a_i c_i & -b_i c_i & a_i^2 + b_i^2 \end{bmatrix} .$$

In Fig. 4, d_i denote the distance from the reconstructed point P to line i.

In principle, 3D seed orientations can also be reconstructed. Given three images, the 3D orientation of a seed can be computed by first defining the three planes that join each projection with its respective X-ray source, and then calculating the intersection line of the three planes. However, due to various errors, these planes never intersect into one straight line in space. So a symbolic intersection becomes necessary. Here we use a simple and quick method to do it (Fig. 5). Here we assume that for a seed S in space, plane i (PL_i) is defined by the projection of S in image i, s_i, with its X-ray source. Since from each pair of planes we can obtain an intersection line, from three planes we can obtain three independent intersection lines. Then we consider the average of the three unit directional vector as the reconstructed 3D orientation of the seed.

Figure 3: The flowchart of MARSHAL

2.4 Cost metrics

The above discussion shows that a robust seed matching algorithm heavily depends on the performance of the cost metric which constructs the cost matrix in MARSHAL. The seed matching from three images has been reduced to a sequence of bipartite matching. For each bipartite matching problem, we construct a cost matrix in the following way. Given three images, we choose two principal images for this bipartite matching problem. For each pair of seeds in these two images, we reconstruct a 3D seed. Then we project this reconstructed 3D seed to the third image, and find the closest seed projection on the image to this projection based on the cost metric. We reconstruct a 3D seed from those three seed projections. Then we project it back to the three images and calculate the average metric value between each projected seed and the observed seed in each image.
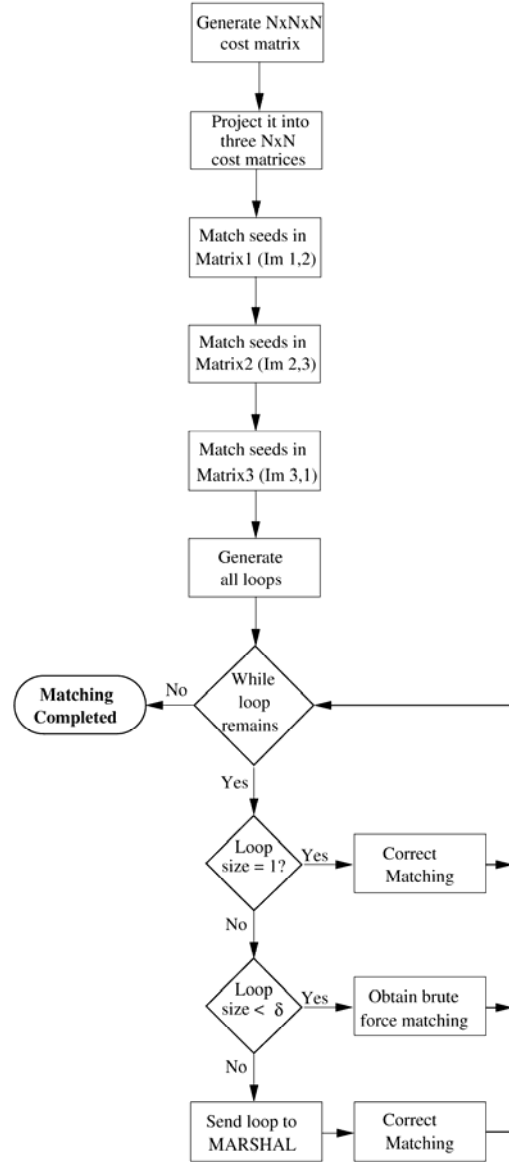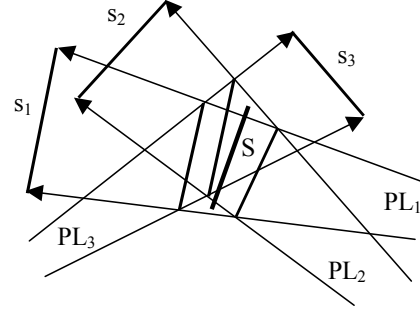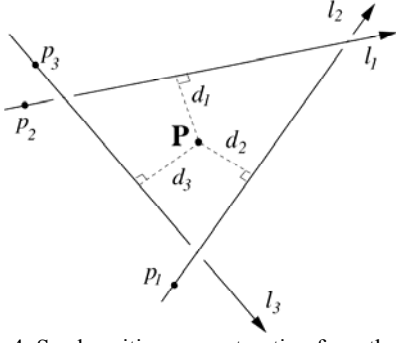
Figure 4: Seed position reconstruction from three images   Figure 5: Seed orientation reconstruction from three images

In our previous work, only the positions of seed centers were used to define the cost metric. However, we believe it is of utmost importance to exploit all the available information in order to minimize the matching and reconstruction errors. Although the information about the orientations and length of seeds may have less significant dosimetric impact, it can positively contribute to improving seed matching rate and 3D implant reconstruction accuracy. It can also become critical information when hidden and spuriously segmented seeds need to be matched, where reliable and generic methods are not yet available. In the following, besides the original cost metric used by MARSHAL (M_1), we will introduce several new cost metrics (M_2- M_7) which integrate seed orientation information. Extensive simulation study has been performed to compare the performance among different cost metrics. The results will be reported in Sec.4. Fig. 6 illustrates the seven metrics tested in this study.

2.4.1 M_1 – cost metric without considering seed orientation

The original MARSHAL uses projection error (PE) to define its cost matrix. PE is obtained by projecting the 3D reconstructed point back onto each image and calculating the mean distance between the projected location and the observed location of the seed. In the original MARSHAL, as seen in Fig. 6, the cost metric M_1 is defined as:

$$M_1 = d_c,$$ (3)

where, s_p and s_o denote the projected seed and the observed seed respectively, a_p and b_p are the end points of s_p, c_p the center of s_p, a_o and b_o the end points of s_o, c_o the center of s_o, and d_c is the center to center distance between the projected seed and the observed seed.
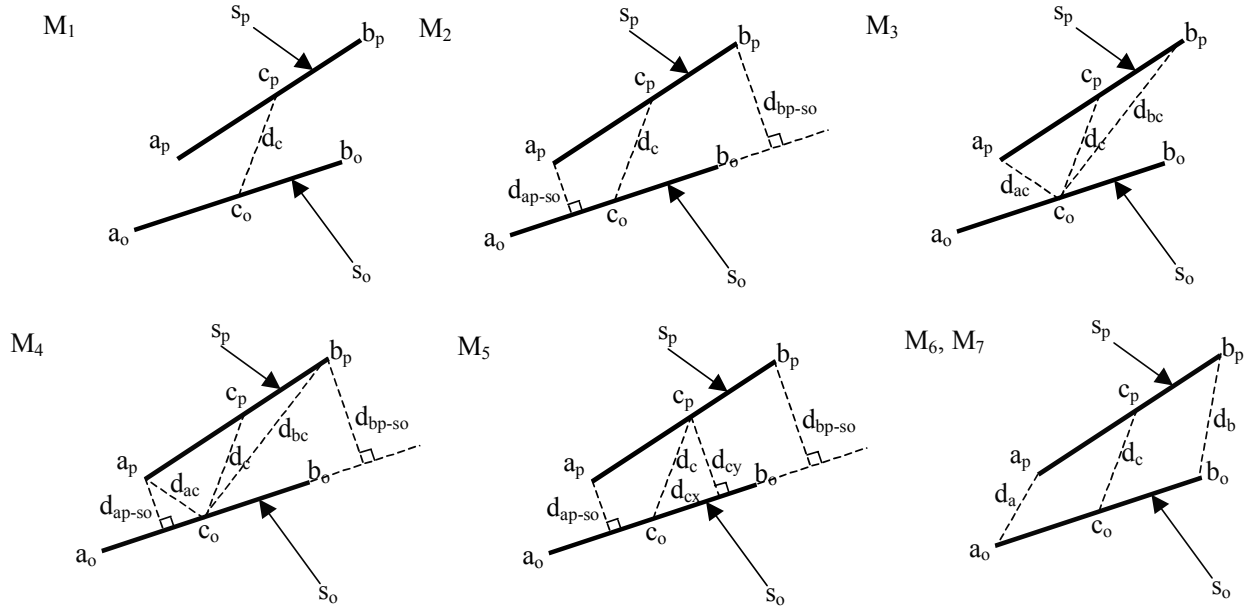

Figure 6: Cost metrics M_1-M_7

Another cost metric (probably the most popular choice) is reconstruction accuracy (RA). RA is obtained by calculating the average distance from the symbolic intersection to each straight line (Fig. 4). It has been proved that PE performs better than RA[59].

2.4.2 $M_2 - M_7$ Cost metrics considering 2D seed orientation
Combining the information of seed orientations with that of seed positions, we have created a series of cost metrics shown in Fig. 6.

M_2 is defined as
$$M_2 = d_c + d_{ap-so} + d_{bp-so},$$
(4)
where d_{ap-so} denotes the distance from a_p to s_o, and d_{bp-so} denotes the distance from b_p to s_o.

M_3 is defined as
$$M_3 = d_c + d_{ac} + d_{bc},$$
(5)
where d_{ac} denotes the distance from a_p to c_o, and d_{bc} denotes the distance from b_p to c_o.

M_4 is defined as
$$M_4 = d_c + d_{ac} + d_{bc} + d_{ap-so} + d_{bp-so}.$$
(6)

M_5 is defined as
$$M_5 = d_{cx} + \frac{d_{cy} + d_{ap-so} + d_{bp-so}}{3},$$
(7)
where d_{cx} is the component of d_c along s_o, and d_{cy} is the component of d_c perpendicular to s_o.

M_6 is defined as
$$M_6 = d_c + d_a + d_b,$$
(8)
where d_a denotes the distance from a_p to a_o, and d_b denotes the distance from b_p to b_o.

M_7 is defined as
$$M_7 = d_a + d_b.$$
(9)

From Fig. 6 and Eq. 4-9, we can see that M_2-M_6 contain not only the center-to-center distance but also the distances related to end points of the seeds. Since the end points define the orientations of the seeds, by including end points into the cost metrics, we have actually incorporated the seed orientation information into the metrics.

2.4.3 M_8- cost metric using threshold for 3D seed orientation
In this case, we define a quantity similar to the RA, the reconstruction accuracy for 3D orientation (RAO). RAO is obtained by calculating the average angular difference from the reconstructed 3D seed orientation to each intersection line obtained from each pair of planes (Fig. 5). In general, for actually matched seeds, the orientations of the three intersection lines are considered not very different from each other. This means that the RAO for actually matched seeds should be small. By setting a threshold for RAO, we define the cost metric M_8 as
$$M_8 = \begin{cases} d_c & RAO \leq threshold \\ \infty & RAO > threshold \end{cases}.$$
(10)
In this way, we eliminate some potential mismatches early at the stage of constructing cost matrices.

Moreover, we notice that for actually matched seeds, the RA and RAO should be small at the same time. Based on this observation, we propose a simple way to find the optimal threshold for RAO. In order to find the optimal threshold, we choose a range of threshold angles. For each threshold angle in the range, we run MARSHAL with M_8 as the cost metric, and calculate the total RA for the resulting matching. Then we pick the threshold angle with the smallest total RA as the optimal threshold.

3. SIMULATION RESULTS AND DISCUSSION

Extensive simulation study has been done on synthetic images to analyze the performance of MARSHAL under various cost metrics. Matlab program was created to model X-ray imaging. Given C-arm parameters and implant details, it generated synthetic images and exact locations and orientations of seeds in the images. In this paper we use three images to evaluate the correspondences with all the above-mentioned cost metrics. To evaluate the performance on simulated data, we compute the percentage of correct matching. We also study the sensitivity of different cost metrics to segmentation errors, because seed segmentation from the X-ray images is one of the most important sources of error. Here we consider the impact of the segmentation errors in both position and orientation. Random error was modeled following uniform distribution. Therefore, a 1 mm position error means that a maximum error of magnitude 1 mm was added to the positions of the seeds following the uniform distribution. Similarly, a 1° orientation error means that a maximum error of magnitude 1° was added to the orientations of the seeds.

Simulation data sets were generated to simulate different seed densities, prostate volumes and image separation angles. Here we focus on large and dense implants. In our simulation study, the seed density varies from 2.0 seeds/cc to 3.0 seeds/cc, the prostate volume from 40cc to 45cc, and image separation angle from 15° to 20°. So the total number of seeds is from 80 to 140. To test the impact of segmentation errors on the cost metrics, we vary the segmentation error in position from 0 mm to 1.5 mm, and the segmentation error in orientation from 0° to 30°. The averaged results are displayed in Tables 1 and 2. The general trend can be seen in Fig. 7 and 8. When we test the metrics with orientation segmentation error, we fix the position segmentation error as 0.75 mm; when we test the metrics with position segmentation error, we fix the orientation error as 9°. This is because they are close to the segmentation errors in practice.

Matching rate (%)			Orientation segmentation errors										
			0°	3°	6°	9°	12°	15°	18°	21°	24°	27°	30°
Cost Metrics	M_1	Avg.	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%
		STD	2.3%	2.3%	2.3%	2.3%	2.3%	2.3%	2.3%	2.3%	2.3%	2.3%	2.3%
	M_2	Avg.	98.4%	98.1%	97.6%	96.2%	93.8%	93.2%	91.8%	91.5%	91.0%	90.3%	90.8%
		STD	1.8%	1.8%	2.2%	2.7%	3.7%	3.0%	3.8%	3.3%	4.0%	4.3%	4.1%
	M_3	Avg.	88.2%	91.5%	93.5%	93.1%	92.4%	92.9%	91.8%	91.8%	92.6%	91.7%	92.0%
		STD	5.7%	3.7%	4.0%	4.0%	3.1%	3.4%	3.6%	4.1%	3.4%	3.9%	3.7%
	M_4	Avg.	96.3%	96.9%	96.4%	95.1%	92.9%	92.8%	90.7%	90.7%	90.6%	89.3%	90.7%
		STD	2.6%	2.2%	2.5%	3.2%	3.8%	2.8%	3.8%	3.9%	3.1%	4.3%	3.6%
	M_5	Avg.	98.1%	98.0%	97.9%	97.1%	95.5%	95.7%	95.1%	94.0%	94.8%	94.0%	94.6%
		STD	1.8%	1.8%	2.1%	2.1%	3.0%	2.9%	3.0%	3.0%	2.8%	3.6%	3.4%
	M_6	Avg.	98.4%	98.2%	98.1%	97.2%	95.4%	94.6%	93.5%	92.2%	93.3%	91.6%	92.8%
		STD	1.6%	1.8%	1.9%	2.2%	3.0%	2.9%	3.5%	3.7%	3.5%	3.8%	3.5%
	M_7	Avg.	98.6%	98.4%	98.0%	96.7%	93.8%	92.5%	90.8%	89.5%	89.7%	87.9%	88.0%
		STD	1.6%	1.7%	2.0%	2.4%	3.8%	3.2%	4.2%	3.6%	3.9%	4.9%	4.7%
	M_8	Avg.	96.5%	96.4%	96.5%	96.6%	96.5%	96.6%	96.6%	96.6%	96.3%	96.4%	96.3%
		STD	2.2%	2.2%	2.2%	2.2%	2.2%	2.2%	2.2%	2.2%	2.2%	2.3%	2.3%

Table 1: Matching rate from various cost metrics under different orientation segmentation errors (with a fixed position segmentation error of 0.75 mm)

Reconstruction error (mm)			0°	3°	6°	9°	12°	15°	18°	21°	24°	27°	30°
Cost Metrics	M_1	All	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
		Mat.	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49
		Mis.	0.65	0.65	0.65	0.65	0.65	0.65	0.65	0.65	0.65	0.65	0.65
	M_2	All	0.50	0.50	0.50	0.52	0.55	0.55	0.55	0.59	0.66	0.62	0.65
		Mat.	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.50	0.51	0.50	0.51
		Mis.	0.50	0.57	0.74	0.98	1.10	1.15	1.17	1.40	1.61	1.43	1.56
	M_3	All	0.76	0.59	0.57	0.60	0.54	0.53	0.58	0.55	0.54	0.56	0.56
		Mat.	0.52	0.50	0.50	0.51	0.49	0.49	0.51	0.49	0.49	0.49	0.49
		Mis.	1.71	1.24	1.08	1.15	1.06	0.96	1.21	1.12	1.05	1.13	1.10
	M_4	All	0.52	0.50	0.51	0.54	0.54	0.54	0.58	0.57	0.62	0.61	0.59
		Mat.	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.50	0.49	0.50
		Mis.	0.92	0.66	0.87	1.13	1.15	1.12	1.31	1.20	1.58	1.45	1.24
	M_5	All	0.50	0.50	0.50	0.50	0.51	0.51	0.52	0.52	0.51	0.53	0.51
		Mat.	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49
		Mis.	0.63	0.52	0.56	0.66	0.76	0.80	0.88	0.86	0.79	1.00	0.69
	M_6	All	0.50	0.50	0.50	0.50	0.51	0.52	0.52	0.55	0.57	0.59	0.55
		Mat.	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.50	0.50	0.50	0.49
		Mis.	0.53	0.49	0.50	0.64	0.78	0.92	0.93	1.11	1.11	1.40	1.09
	M_7	All	0.50	0.50	0.50	0.51	0.53	0.56	0.56	0.60	0.65	0.70	0.72
		Mat.	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.50	0.50	0.50	0.51
		Mis.	0.52	0.48	0.58	0.74	0.91	1.27	1.28	1.30	1.58	1.76	1.94
	M_8	All	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
		Mat.	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49	0.49
		Mis.	0.65	0.65	0.64	0.64	0.65	0.64	0.64	0.65	0.72	0.65	0.65

Table 2: Average reconstruction error from various cost metrics under different orientation segmentation errors (with a fixed position segmentation error of 0.75 mm). "Mat." stands for matched seeds, "Mis." stands for mismatched seeds.

Matching rate (%)			0	0.25	0.5	0.75	1.0	1.25	1.5
Cost Metrics	M_1	Avg.	100.0%	99.5%	98.5%	96.0%	92.5%	86.7%	76.2%
		STD	0.0%	0.9%	1.7%	2.2%	4.0%	4.9%	6.1%
	M_2	Avg.	99.3%	98.8%	98.1%	96.6%	93.6%	87.5%	78.4%
		STD	0.9%	1.2%	1.9%	2.4%	3.8%	4.9%	7.7%
	M_3	Avg.	98.5%	97.4%	95.7%	93.5%	90.2%	83.3%	73.6%
		STD	1.4%	2.0%	2.4%	3.5%	4.0%	6.6%	7.9%
	M_4	Avg.	99.0%	98.5%	97.5%	95.8%	92.5%	87.1%	77.6%
		STD	1.3%	1.6%	1.7%	3.1%	3.7%	6.0%	7.4%
	M_5	Avg.	99.8%	99.4%	98.6%	97.1%	94.2%	89.6%	80.7%
		STD	0.3%	1.0%	1.6%	2.2%	3.6%	4.7%	6.3%
	M_6	Avg.	99.6%	99.0%	98.5%	97.2%	95.1%	90.1%	82.1%
		STD	0.7%	1.1%	1.6%	2.4%	3.1%	5.2%	6.4%
	M_7	Avg.	98.8%	98.6%	97.9%	96.8%	94.8%	90.7%	83.1%
		STD	1.3%	1.3%	1.9%	2.4%	3.3%	5.2%	6.2%
	M_8	Avg.	100.0%	99.5%	98.6%	96.2%	92.9%	88.2%	78.3%
		STD	0.0%	0.9%	1.5%	2.3%	3.9%	4.9%	5.6%

Table 3: Matching rate of various cost metrics under different position segmentation errors (with a fixed orientation segmentation error of 9°)

Reconstruction error (mm)			Position segmentation errors (mm)						
			0	0.25	0.5	0.75	1.0	1.25	1.5
Cost Metrics	M_1	All	0.00	0.16	0.32	0.50	0.71	1.21	2.01
		Mat.	0.00	0.16	0.32	0.48	0.63	0.84	1.06
		Mis.	0.00	0.08	0.38	0.82	1.44	2.94	4.06
	M_2	All	0.01	0.17	0.32	0.50	0.72	1.20	2.07
		Mat.	0.00	0.16	0.32	0.48	0.63	0.83	1.07
		Mis.	0.14	0.25	0.39	0.80	1.61	3.26	4.33
	M_3	All	0.01	0.17	0.34	0.53	0.82	1.48	2.32
		Mat.	0.00	0.16	0.32	0.48	0.66	0.85	1.13
		Mis.	0.29	0.45	0.62	1.06	1.85	3.65	4.34
	M_4	All	0.01	0.17	0.34	0.51	0.74	1.32	2.09
		Mat.	0.00	0.16	0.32	0.48	0.64	0.84	1.09
		Mis.	0.26	0.33	0.73	0.91	1.58	3.45	4.23
	M_5	All	0.00	0.16	0.32	0.50	0.69	1.10	1.78
		Mat.	0.00	0.16	0.32	0.48	0.63	0.81	1.04
		Mis.	0.04	0.10	0.36	0.77	1.30	2.77	3.63
	M_6	All	0.00	0.16	0.32	0.49	0.67	1.07	1.69
		Mat.	0.00	0.16	0.31	0.48	0.63	0.81	1.01
		Mis.	0.09	0.20	0.38	0.69	1.22	2.49	3.50
	M_7	All	0.01	0.17	0.33	0.50	0.67	1.00	1.58
		Mat.	0.00	0.16	0.32	0.48	0.63	0.80	1.02
		Mis.	0.33	0.36	0.56	0.87	1.28	2.19	3.13
	M_8	All	0.00	0.16	0.32	0.50	0.70	1.08	1.79
		Mat.	0.00	0.16	0.32	0.48	0.63	0.82	1.05
		Mis.	0.00	0.07	0.38	0.82	1.43	2.42	3.53

Table 4: Average reconstruction error from various cost metrics under different position segmentation errors (with a fixed orientation segmentation error of 9°). "Mat." stands for matched seeds, "Mis." stands for mismatched seeds.

From the tables and figures, we can see two clear trends:

1. When the position segmentation error increases, the matching rate based on each cost metric decreases monotonically. This is because all the cost metrics we discussed depend on the distance between the projected seeds and the observed seeds. The increase in position segmentation error will cause the difference between the metric values of actually matched seeds and those of actually unmatched seeds decrease.

2. When the orientation segmentation error increase, the matching rate of the cost metrics considering orientation decreases. This is because these cost metrics depend on reasonably accurate 2D orientation information. The increase in orientation segmentation error will cause the difference between the metric values of actually matched seed and those of actually unmatched seeds decrease.

Moreover, if we use M_1 (the cost metric without considering seed orientation) as a criterion, M_8 (the cost metric using threshold for 3D seed orientation) always give us a matching rate equal to or slightly greater than that of M_1. This is because M_8 is built upon M_1 with a threshold. It guarantees that M_8 performs at least as well as M_1. However, it also means that no substantial improvement over M_1 can be obtained from M_8. The data show that the matching rate of M_3 is always below that of M_1. The matching rate of M_4 drops below that of M_1 right after the orientation segmentation error hits 6°, while its sensitivity to the position segmentation error is about same as M_1. Therefore, M_3, M_4 and M_8 are not recommended.

The data also show that the matching rates of M_2, M_5, M_6 and M_7 are higher than that of M_1 when the orientation segmentation error is lower than 9°. Among them M_5 and M_6 are better than M_2 and M_7. When the orientation segmentation error is equal to 9°, M_5 and M_6 give the matching rates greater than 97% while M_2 and M_7 are close to M_1, about 96%. Meanwhile, the standard deviation of the matching rates of M_5 and M_6 are also slightly smaller than that of M_1. M_5 and M_6 also give the best matching rates when the position segmentation error is close to the practice (0.75mm). At the same time, the average reconstruction errors of mismatched seeds from M_5 and M_6 are about same as or slightly

lower than that of M_1 at the segmentation errors of 9° and 0.75mm. Therefore M_5 and M_6 are the best cost metrics we have. They can be used to robustly match and reconstruct seeds under normal segmentation errors. It also means that combining seed orientation information into the cost metric does result in certain improvement in the performance of the algorithm.
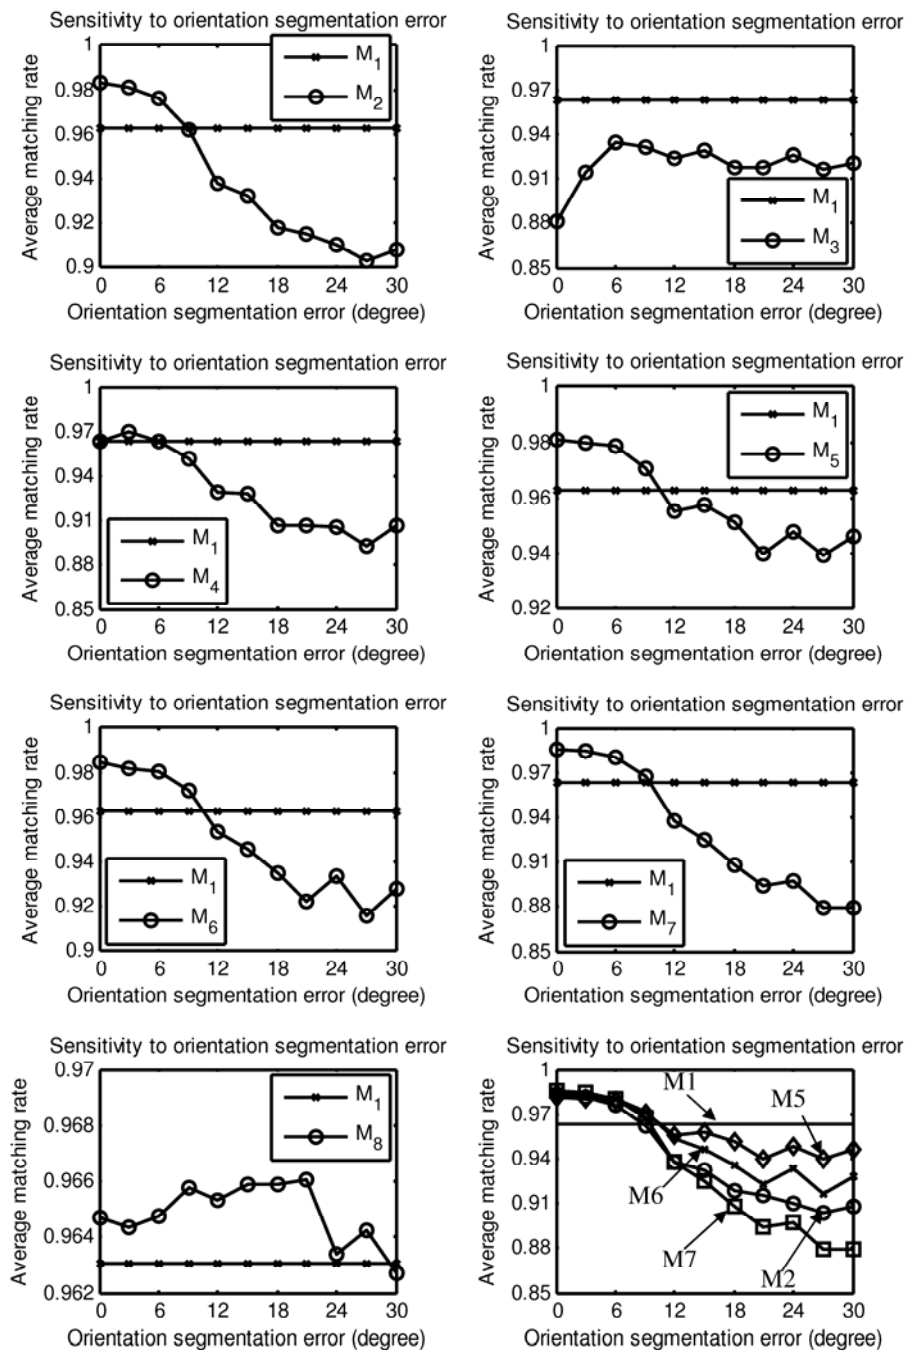

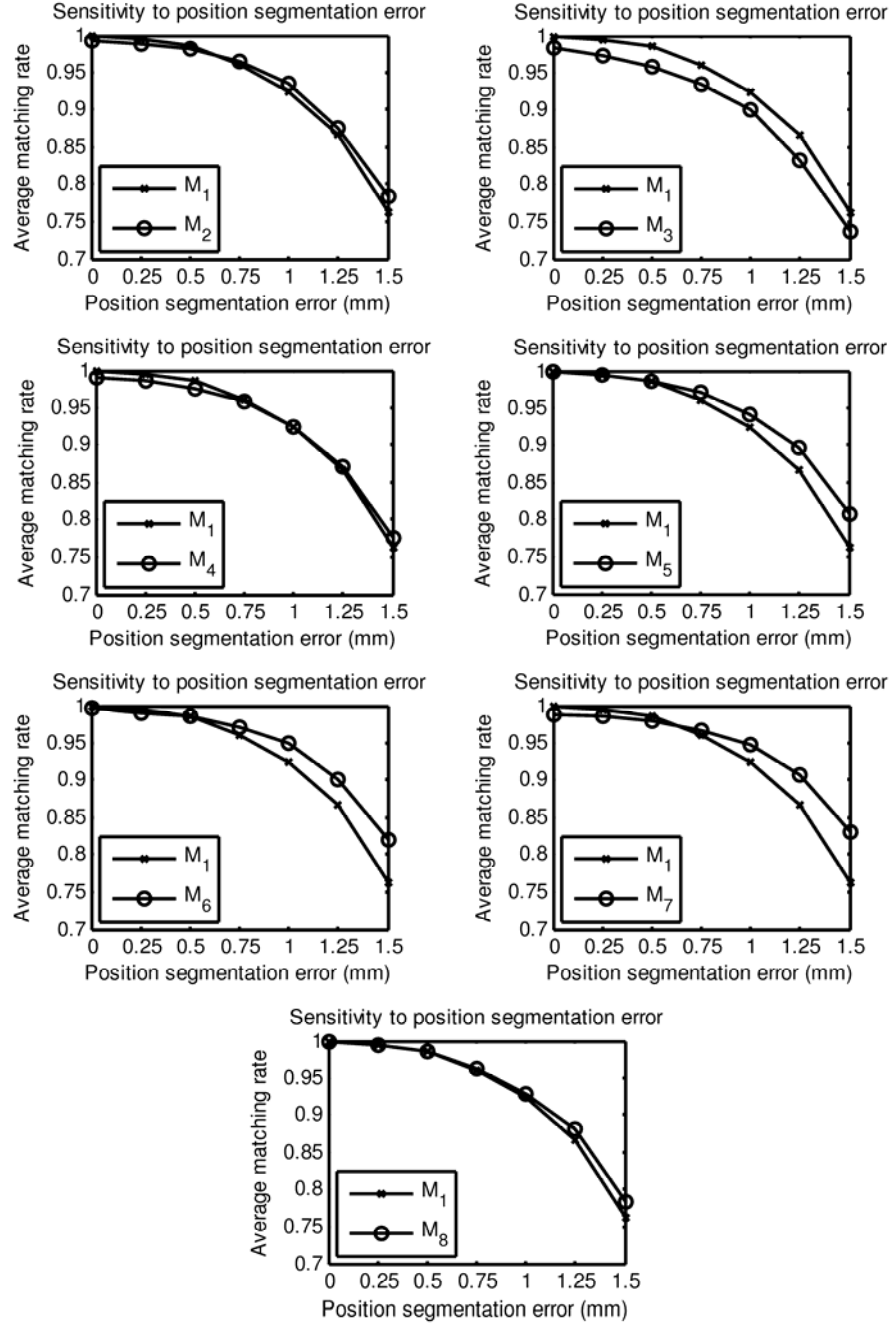Figure 7: Sensitivity of cost metrics to orientation segmentation error

Figure 8: Sensitivity of cost metrics to position segmentation error

4. CONCLUSIONS AND FUTURE WORK

In this paper, we have extended our previously proposed quasi-polynomial time algorithm (MARSHAL) to combine the information of seed orientations into matching procedure. Simulation results show that the cost metrics considering seed orientation can further improve the matching rate of MARSHAL, and robust to certain amount of seed segmentation errors in both position and orientation.

An underlying assumption of MARSHAL is that all the seeds are segmented and their 2D positions and orientations are known. In reality, however, some seeds always remain hidden, and some are segmented spuriously[35,63,64]. In a parallel work[63], we extended MARSHAL to deal with hidden seeds. In phantom study, we managed to recover >96% seeds from three images and >99% from four images, and the algorithm was robust to segmentation error up to 1 mm and hidden seed rate up to 8%. While the matching rate was excellent, the reconstruction error for mismatched seeds remained 3-8mm. In this paper, Sec. 3 shows that under average segmentation errors, cost metrics that utilize seed orientation significantly lower the reconstruction error for mismatched seeds, generally to less than 1mm.

We conjecture that combining these new cost metrics in our hidden seed recovery method[63] will result in superior seed matching rate and low reconstruction error for both matched and unmatched seeds, and thereby yield a clinically applicable solution.

ACKNOWLEDGEMENTS

REFERENCES

1. A. Jemal, "Cancer statistics," Ca-Cancer J. Clin. **54**(1), 8–29, 2004.
2. R. Peschel and J. Colberg, "Surgery, brachytherapy, and external-beam radiotherapy for early prostate cancer," Lancet Oncol. **4**, 233–241, 2003.
3. J. Blasko, T. Mate, J. Sylvester, P. Grimm, and W. Cavanagh, "Brachytherapy for carcinoma of the prostate: Techniques, patient selection, and clinical outcomes," Semin. Radiat. Oncol. **12**, 81–94, 2002.
4. G. Merrick, W. Butler, J. Lief, and A. Dorsey, "Is brachytherapy comparable with radical prostatectomy and external-beam radiation for clinically localized prostate cancer?," Tech. Urol. **7**, 12–19, 2001.
5. S. Nag, J. Ciezki, R. Cormack, S. Doggett, K. DeWyngaert, G. Edmundson, R. Stock, N. Stone, Y. Yu, and M. Zelefsky, "Intraoperative planning and evaluation of permanent prostate brachytherapy: Report of the American Brachytherapy Society," Int. J. Radiat. Oncol., Biol., Phys. **51**, 1422–1430, 2001.
6. J. Xue, E. Gressen, and T. Jefferson, "Feasibility of trus-based prostate post-implant dosimetry," AAPM Annual Meeting, July 2004, Poster.
7. V. A. Dumane, M. Zaider, G. N. Cohen, and W. M. F. Worman, "Combined ultrasound-fluoroscopy approach to the intraoperative detection of seeds in prostate brachytherapy," ASTRO Annual Meeting, Poster, Oct 3–7 2004.
8. F. Mitri, P. Trompette, and J. Chapelon, "Improving the use of vibroacoustography for brachytherapy metal seed imaging: A feasibility study," IEEE Trans. Med. Imaging **23**, 1–6, 2004.
9. D. Holmes, B. Davis, C. Bruce, and R. Robb, "3d visualization, analysis, and treatment of the prostate using trans-urethral ultrasound," Comput. Med. Imaging Graph. **27**, 339–349, 2003.
10. D. French, J. Morris, M. Keyes, and S. E. Salcudean, "Real-time dosimetry for prostate brachytherapy using trus and fluoroscopy," MICCAI 2004, 983–991.
11. B. Han, K. Wallner, G. Merrick, W. Butler, S. Sutlief, and J. Sylvester, "Prostate brachytherapy seed identification on post-implant trus images," Med. Phys. **30**, 898–900, 2003.
12. P. Kumar, R. Good, B. Epstein, M. Hussain, and F. Bartone, "Fluoroscopy guided transperineal percutaneous permanent 125iodine implantation of prostate cancer," Radiat. Med. **3**, 161–167, 1985.
13. B. Prestidge, J. Prete, T. Buchholz, J. Friedland, R. Stock, P. Grimm, and W. Bice, "A survey of current clinical practice of permanent prostate brachytherapy in the United States," Int. J. Radiat. Oncol., Biol., Phys. **15**, 461–465, 1998.
14. P. Ravindran, C. Lewis, J. Alphonsi, P. Lindsay, and D. D'Souza, "A c-arm based intra-operative dosimetry system for trans-rectal ultrasound guided prostate implant with 125I seeds," in Proceedings of the 14th International Conference on the Use of Computers in Radiation Therapy, 2004.
15. D. Todor, M. Zaider, G. Cohen, M. Worman, and M. Zelefsky, "Intraoperative dynamic dosimetry for prostate implants," Phys. Med. Biol. **48**, 1153–1171, 2003.
16. D. Todor, G. Cohen, H. Amols, and M. Zaider, "Operator-free, film-based 3d seed reconstruction in brachytherapy," Phys. Med. Biol. **47**, 2031–2048, 2002.
17. L. Gong, P. Cho, B. Han, K. Wallner, S. Sutlief, S. Pathak, D. Haynor, and Y. Kim, "Ultrasonography and fluoroscopic fusion for prostate brachytherapy dosimetry," Int. J. Radiat. Oncol., Biol., Phys. **1**, 1322– 1330, 2002.

18. G. Grado, T. Larson, C. Balch, M. Grado, J. Collins, J. Kriegshauser, G. Swanson, R. Navickis, and M. Wilkes, "Actuarial disease-free survival after prostate cancer brachytherapy using interactive techniques with biplane ultrasound and fluoroscopic guidance," Int. J. Radiat. Oncol., Biol., Phys. **1**, 289–299, 1998.

19. G. Merrick, W. Butler, A. Dorsey, and H. Walbert, "Prostatic conformal brachytherapy: 125i/103pd postoperative dosimetric analysis," Radiat. Oncol. Invest. **5**, 305–313, 1997.

20. D. Nori and J. Moni, "Current issues in techniques of prostate brachytherapy," Semin Surg. Oncol. **13**, 444–453, 1997.

21. S. Nag, D. Scaperoth, R. Badalament, S. Hall, and J. Burgers, "Transperineal palladium 103 prostate brachytherapy: Analysis of morbidity and seed migration," Urology **45**, 87–92, 1995.

22. K. Wallner, S. Chiu-Tsao, J. Roy, V. Arterbery, W. Whitmore, S. Jain, B. Minsky, P. Russo, and Z. Fuks, "An improved method for computerized tomography-planned transperineal 125iodine prostate implants," J. Urol. **146**, 90–95, 1991.

23. R. Hofstetter, M. Slomczykowski, M. Sati, and L. Nolte, "Fluoroscopy as an imaging means for computer-assisted surgical navigation," Comput. Aided Surg. **4**, 65–76, 1999.

24. J. Yao, R. H. Taylor, R. P. Goldberg, R. Kumar, A. Bzostek, V. R. Van, P. Kazanzides, and A. Gueziec, "A c-arm fluoroscopy-guided progressive cut refinement strategy using a surgical robot," Comput. Aided Surg. **5**, 373–390, 2000.

25. A. Jain, T. Mustufa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (ftrac) fiducial," Med. Phys. **32**, 3185–3198, 2005.

26. A. Fung, "C-arm imaging for brachytherapy source reconstruction: Geometrical accuracy," Med. Phys. **29**, 724–726, 2002.

27. S. Narayanan and P. Cho, "3d seed reconstruction from unknown imaging geometry," in Proceedings of the 14th International Conference on the Use of Computers in Radiation Therapy, 2004, 762–765.

28. M. Murphy and D. Todor, "Demonstration of a forward iterative method to reconstruct brachytherapy seed configurations from x-ray projections," Phys. Med. Biol. **50**, 2715–2737, 2005.

29. A. Jain, R. Kyon, Y. Zhou and G. Fichtinger, "C-arm calibration—Is it really necessary?," Medical Image Computing and Computer Assisted Interventions (MICCAI), 2005.

30. K. Wallner, J. Roy, M. Zelefsky, Z. Fuks, and L. Harrison, "Fluoroscopic visualization of the prostatic urethra to guide transperineal prostate implantation," Int. J. Radiat. Oncol., Biol., Phys. **1**, 863–867, 1994.

31. P. Roberson, V. Narayana, D. McShan, R. Winfield, and P. McLaughlin, "Source placement error for permanent implant of the prostate," Med. Phys. **24**, 251–257, 1997.

32. M. Baird, R. Holt, and T. Selby, "Improvement of transperineal implant dosimetry by intraoperative cystoscopic confirmation of prostate anatomy," J. Urol. **164**, 406–410, 2000.

33. L. Archambault, L. Beaulieu, and D. Tubic, "Automatic post-implant needle reconstruction algorithm to characterize and improve implant robustness analyses," Med. Phys. **30**, 2897–2903, 2003.

34. M. Zhang, M. Zaider, M. Worman, and G. Cohen, "On the question of 3d seed reconstruction in prostate brachytherapy: The determination of x-ray source and film locations," Phys. Med. Biol. **49**, N335–345, 2004.

35. Y. Su, B. Davis, M. Herman, and R. Ra, "Prostate brachytherapy seed localization by analysis of multiple projections: Identifying and addressing the seed overlap problem," Med. Phys. **31**, 1277–1287, 2004.

36. D. Tubic, A. Zaccarin, J. Pouliot, and L. Beaulieu, "Automated seed detection and three-dimensional reconstruction. II. Reconstruction of permanent prostate implants using simulated annealing," Med. Phys. **28**, 2272–2279, 2001.

37. S. Lam, P. Cho, and R. Marks, "Prostate brachytherapy seed segmentation using spoke transform," Proc. SPIE **4322**, 1490–1500, 2001.

38. P. Cho, "Computerized segmentation of clustered seeds in prostate brachytherapy," International Conference on the Use of Computers in Radiation Therapy, 2000, 105–107.

39. H. Amols and I. Rosen, "A three-film technique for reconstruction of radioactive seed implants," Med. Phys. **8**, 210–214, 1981.

40. P. Biggs and D. Kelley, "Geometric reconstruction of seed implants using a three-film technique," Med. Phys. **10**, 701–704, 1983.

41. M. Rosenthal and R. Nath, "An automatic seed identification technique for interstitial implants using three isocentric radiographs," Med. Phys. **10**, 475–479, 1983.

42. M. Altschuler, P. Findlay, and R. Epperson, "Rapid accurate, threedimensional location of multiple sees in implant radiotherapy treatment planning," Phys. Med. Biol. **28**, 1305–1308, 1983.

43. D. Jackson, "An automatic method for localizing radioactive seeds in implant dosimetry," Med. Phys. **10**, 370–372, 1983.

44. S. Li, G. Chen, C. Pelizzari, C. Reft, J. Roeske, and Y. Lu, "A new source localization algorithm with no requirement of one-to-one source correspondence between biplane radiographs," Med. Phys. **23**, 921–927, 1996.

45. R. L. Siddon and L. M. Chin, "Two-film brachytherapy reconstruction algorithm," Med. Phys. **12**, 77–83, 1985.

46. M. Altschuler and A. Kassaee, "Automated matching of corresponding seed images of three simulator radiographs to allow 3d triangulation of implanted seeds," Phys. Med. Biol. **42**, 293–302, 1997.

47. S. Narayanan, P. Cho, and R. Marks, "Fast cross-projection algorithm for reconstruction of seeds in prostate brachytherapy," Med. Phys. **29**, 1572–1579, 2002.

48. K. Thornton, "Prostate brachytherapy seed localization by fluoroscopy/ ultrasound fusion: Algorithms and analysis," AAPM Annual Meeting, July 2004.

49. H. Liu, G. Cheng, Y. Yu, R. Brasacchio, D. Rubens, J. Strang, L. Liao, and E. Messing, "Automatic localization of implanted seeds from postimplant ct images," Phys. Med. Biol. **48**, 1191–1203, 2003.

50. D. Dubois, W. J. Bice, and B. Prestige, "Ct and mri derived source localization error in a custom prostate phantom using automated image coregistration," Med. Phys. **28**, 2280–2284, 2001.

51. K. Shafique and M. Shah, "A non-iterative greedy algorithm for multiframe point correspondence," ICCV, 2003.

52. C. J. Veenman, M. J. T. Reinders, and E. Backer, "Resolving motion correspondence for densely moving points," IEEE Trans. Pattern Anal. Mach. Intell. **23**, 54–72, 2001.

53. C. Veenman, M. Reinders, and E. Backer, "Establishing motion correspondence using extended temporal scope," Artificial Intelligience **145**, 227–243, 2003.

54. V. Salari and I. Sethi, "Feature point correspondence in the presence of occlusion," Pattern Anal. Machine Intell. **12**, 87–91, 1990.

55. Y. Cheng, R. Collins, A. Hanson, and E. Riseman, "Triangulation without correspondences," ARPA94, 1994, II:993–1000.

56. Y. Cheng, V. Wu, R. Collins, A. Hanson, and E. Riseman, "Maximumweight bipartite matching technique and its application in image feature matching," 1996.

57. *Artificial Intelligence: A Modern Approach*, Prentice Hall Series in Artificial Intelligence, edited by S. Russell and P. Norvig, 2nd ed..

58. T. H. Cormen, C. E. Leiserson, and R. Rivest, *Introduction to Algorithms* 2nd ed. MIT Press, 1990.

59. A. K. Jain, Y. Zhou, T. Mustufa, E. C. Burdette, G. S. Chirikjian, G. Fichtinger, "Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)", Med. Phys. **32** (11), 3475-3492, 2005.

60. R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, 1st ed. Prentice Hall, Englewood Cliffs, NJ, 1993.

61. H. W. Kuhn, "The Hungarian method for the assignment problem," Naval Res. Logistics Quart. **2**, 83–97, 1955.

62. C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity,* Prentice-Hall, Englewood Cliffs, NJ, 1982.

63. R. C. Kon, A. Jain and G. Fichtinger, "Hidden seed reconstruction from C-arm images in brachytherapy", ISBI 2006, Washington D.C., April 2006 (accepted).

64. S. Narayanan, P. S. Cho and R. J. Marks, "Three-dimensional seed reconstruction from an incomplete data set for prostate brachytherapy", Phys. Med. Biol. **49** (15): 3483-3494, 2004.

C-arm Calibration - Is it Really Necessary?⋆

Ameet Jain[1], Ryan Kon[1], Yu Zhou[2], and Gabor Fichtinger[1,2]

[1] Department of Computer Science, Johns Hopkins University
[2] Department of Mechanical Engineering, Johns Hopkins University

Abstract. C-arm fluoroscopy is modelled as a perspective projection, the parameters of which are estimated through a calibration procedure. It has been universally accepted that precise intra-procedural calibration is a prerequisite for accurate quantitative C-arm fluoroscopy guidance. Calibration, however, significantly adds to system complexity, which is a major impediment to clinical practice. We challenge the status quo by questioning the assumption that precise intra-procedural calibration is really necessary. We derived theoretical bounds for the sensitivity of 3D measurements to mis-calibration. Experimental results corroborated the theory in that mis-calibration in the focal spot by as much as 50 mm still allows for tracking with an accuracy of 0.5 mm in translation and 0.65^o in rotation, and such mis-calibration does not impose any additional error on the reconstruction of small objects.

1 Introduction

C-arm fluoroscopy is ubiquitous in general surgery, due to its real-time nature, versatility, and low cost. At the same time, quantitative fluoroscopy has not found a *large scale clinical acceptance*, because of inherent technical difficulties and needs to solve four major problems: (1) C-arm image distortion; (2) Calibration of model parameters; (3) Pose recovery or tracking when multiple images are taken; and (4) Registration to imaging modalities. Some of the prominent works that have tackled the above problems are [1, 2]. The driving application of our research is prostate brachytherapy, where radioactive seeds are required to be precisely placed into the prostate. Quantitative fluoroscopy could enable a significant improvement in the current clinical practice.

If is known that both image distortion[3] and calibration[4] may vary significantly with pose. Image distortion usually has a consequential contribution to reconstruction error and needs to be compensated. Thus the additional cost of a full online calibration is not substantial. Recently developed advanced intensifier tubes allow for lesser distortion, while modern flat panel detectors obviate distortion correction altogether. This fact brings up the question whether we need to calibrate the C-arm fully at each pose. The question also leads to the broader issue, that even if it is not pose dependent, how accurate does calibration need to be. In spite of the importance of calibration in C-arm fluoroscopy, as far as the authors are aware, there has been no prior work that conducts this analysis. The

---

vision community has a similar problem [5, 6] when cameras are used for visual serving of robots. We do not go into a detailed comparison for lack of space.

In quantitative C-arm fluoroscopy, we typically need to measure the spatial transformation between two objects, such as a vertebra and a bone drill, as compared to the transformation between an object and the C-arm itself. Thus the central intuition of this paper is that *while an incorrect calibration gives erroneous estimates for the absolute transformations, nevertheless it still provides acceptable relative estimates.* The consequence of this conjecture is potentially far reaching, as it can turn fluoroscopy to an affordable quantitative measurement tool in a large family of procedures. It should be however noted that we do not claim that calibration would always be unnecessary, since there are many applications that require high reconstruction accuracy. The decision should always be made case by case, experimentally. In this paper, we build a mathematical framework to formally address this issue and lend credit to the intuition that a loose estimate of the C-arm parameters might suffice in applications where the relative pose of objects is to be measured. In particular, we prove in theory and demonstrate experimentally that intra-operative dosimetry of brachytherapy implants is feasible with an un-calibrated C-arm.

2   Mathematical Framework

C-arm Imaging Model: Geometric aspects of fluoroscopic imaging can be approximated as a perspective projection with five parameters[7]. There are a total of five independent parameters that need to be evaluated by the calibration procedure - the pixel sizes (two) and the focal spot (three). The pixel sizes are fixed and remain unchanged throughout the life of the C-arm, reducing online-calibration to just the focal spot. Though our framework can study sensitivity due to any of the five parameters, we limit ourselves only to the focal spot.

2.1   Model for Reconstruction Space Transformation

As illustrated in Figure 1(a), let A & B (with reference frames F_A & F_B) be the two objects being imaged. The assumptions are: (i) $^I F_A, ^I F_B$ can be computed from the images; (ii) A & B are *not large* in comparison to the focal length; (iii) F_A and F_B are *close by*; and (iv) the quantity of interest is $^A F_B = (^I F_A)^{-1} \, ^I F_B$. Let $\bar{f}_1$ be the true focal spot and $\bar{f}_2 = (\bar{f}_1 + \bar{D})$ be the mis-calibrated estimate. We claim that even though the absolute locations of the objects are off, their relative transformation might still be accurate.

A transformation is needed that can take the absolute location of an object reconstructed with calibration $\bar{f}_1$, and compute its corresponding location with calibration $\bar{f}_2$. We claim that the simplest transformation will be a linear affine model $\mathcal{T}$. The intuition derives from the observation that the image plane is the same in both reconstruction spaces. Thus if P_1 (not in homogenous coordinates) projects to a point p on the image, then it is constrained to be on line $\bar{L}_1$ in the $\bar{f}_1$-space and on $\bar{L}_2$ in $\bar{f}_2$-space. Thus we seek a continuous invertible
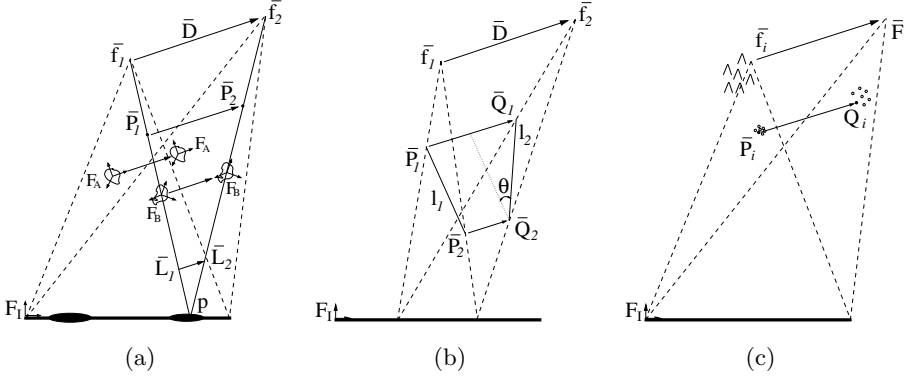
Fig. 1. Mis-calibration (a) shifts all reconstructed objects under an affine transformation; (b) rotates and scales a straight line segment; (c) Pose dependent calibration might be successfully approximated by using the mean value

transformation that projects $\bar{L}_1$ to $\bar{L}_2$. By incorporating the above constraints, $\mathcal{T}$ can be evaluated to be,

$$
P_2 \;=\; \mathcal{T} \cdot P_1 \;=\; \begin{bmatrix} 1 & 0 & D_x/f_{1z} \\ 0 & 1 & D_y/f_{1z} \\ 0 & 0 & 1+(D_z/f_{1z}) \end{bmatrix} \cdot P_1 \;=\; P_1 + (d \cdot Z/f_{1z})\hat{D} \quad (1)
$$

where with respect to (wrt) F_I, $\bar{D} = (D_x, D_y, D_z)$; $d = \|\bar{D}\|_2$; $\hat{D} = \bar{D}/d$; $\bar{f}_1 = (f_{1x}, f_{1y}, f_{1z})$; and $P_1 = (X, Y, Z)$. Each point is effectively translated in direction $\hat{D}$ by an amount proportional to its distance from the image. Experiments measuring the correctness of this affine model are available in Section 3. Thus to study sensitivity, it is sufficient to study the properties of $\mathcal{T}$.

2.2   Changes in Length and Scale

$\mathcal{T}$ preserves the scale along the x, y-axes, but scales the space along the z-axis. Let $P_1(X_1, Y_1, Z_1)$ & $P_2(X_2, Y_2, Z_2)$ be any two points (not necessarily close to each other) in the $\bar{f}_1$-space at a distance of l_1. $\mathcal{T}$ maps them to points Q_1 & Q_2 in the $\bar{f}_2$-space at a distance of l_2 (Figure 1 (b)). It can be shown that

$$
\|l_2 - l_1\| \;\leq\; \frac{d}{f_{1z}}|Z_1 - Z_2| \quad (2)
$$

It directly follows from Equation (2) that $\mathcal{T}$ does not alter the length significantly. As an example, a 10 mm calibration error would affect the length of a 30 mm thoracic pedicle screw at an angle of 45^o by less than 0.2 mm (focal length $\sim$ 1 m), which is significantly less than the error from other sources. Thus F_A, F_B will not change their relative translation by a factor more than that specified by Equation (2).

2.3   Changes in Absolute Orientation

A change in orientation results from the object having a depth (Figure 1 (c)). It can be shown geometrically that the orientation error is maximal when the vector $\overline{P_1 P_2}$ is roughly orthogonal to $\bar{D}$ and is purely in the vertical plane. The amount (θ) and the axis ($\hat{\kappa}$) of rotation, through a series of computations can be shown to be as in Equation (3). The bound on the rotation error is dependent only on origin mis-calibration and not on that in focal length . More importantly it is independent of the height/depth of the object (as far as it is non-planar) and its distance from the image plane. Thus F_A, F_B in Figure 1 will observe the same absolute rotation, in effect not experiencing any relative rotation. Experimental results corroborating this claim are available in Section 3.

$$|\theta| \leq \arcsin \left[\frac{\sqrt{D_x^2 + D_y^2}}{f_{1z}} \right] \sim \frac{\sqrt{D_x^2 + D_y^2}}{f_{1z}} \; ; \quad \hat{\kappa} = \frac{1}{\sqrt{D_x^2 + D_y^2}} (D_y, -D_x, \; 0) \quad (3)$$

2.4   Error in Reconstruction of Point Features

In many applications (particularly in brachytherapy), C-arms are used to reconstruct 3D point objects. This is done by obtaining multiple images at varying orientations and then using triangulation to obtain the desired intersection. In ideal circumstances, all the lines would intersect at a unique point. In practice however, calibration (and other) errors lead to non-intersecting lines. We will attempt to bound the error in this symbolic reconstruction of the point. Let point P be imaged from N different poses and reconstructed in a tracker frame F_A, which is stationary wrt P. Let the i^{th} pose have a focal spot error (in frame F_A) of $\bar{D}_i$. Without errors, each reconstructed line (l_i) would pass through P. It can be shown that due to the calibration error $\bar{D}_i$, the new line passes through a new point $\bar{P}'_A$ and undergoes a rotation ϕ.

$$\bar{P}'_A \sim \bar{P}_A + [0 \;\; 0 \;\; \frac{(\bar{P}_A \cdot \bar{D}_i)}{f_{iz}}]' \quad ; \qquad \phi \sim \frac{(\hat{l}_i \cdot \bar{D}_i) \; sin\theta_i}{f_{iz}} \qquad (4)$$

where θ_i is the angle that l_i makes with the z-axis of F_A. The rotation is fairly small and can be ignored. Thus P_A is at a distance of $(\bar{P}_A \cdot \bar{D}_i) sin\theta_i / f_{iz}$ from l_i. If Q is the symbolic intersection of all l_i's, then it can be shown that Q is no further away than $(\frac{d_{max}}{f_z} sin\theta_{max}) \|P_A\|$ away from any of the lines. Moreover, the reconstruction error (RE) can also be shown to be bounded by

$$RE = \|(\bar{Q} - \bar{P}_A)\| < \frac{\sqrt{2} \, d_{max}}{f_z} \, \|\bar{P}_A\| \qquad (5)$$

where d_{max} is the maximum amount of mis-calibration and f_z is the minimum focal length. Thus a 10 mm focal length error causes an error less than 0.5 mm for a point at a distance of 35 mm. Note that this is the worst case error analysis and in practice the dot product in Equation (4) mutually cancels positive and negative errors, leading to extremely low reconstruction errors (Section 3).

2.5   The Optimal Choice for Calibration Parameters

Since the focal spot is pose dependant, and the results from Section 2.2, & 2.3 suggest robustness to mis-calibration, choosing a constant calibration for quantitative reconstruction might be viable. In the scenario that the focal spot might vary as much as 10 mm from one pose to another, "*what constant calibration should be chosen to minimize error*"?

Let us assume that we are imaging a point P from N different poses (Figure 1 (c)). Wrt frame F_I, let the i^{th} pose have the focal spot at $\bar{f}_i = (f_{ix}, f_{iy}, f_{iz})$ and the point be at location $P_i = (X_i, Y_i, Z_i)$. Note that we assume: (a) variations in each of $f_{ix}, f_{iy}, f_{iz}, X_i, Y_i, Z_i$ & pose are independent; (b) P_i's are close to the iso-center, *i.e.* variations in X_i, Y_i, Z_i are not high. We choose a constant value of $\bar{F} = (F_x, F_y, F_z)$ for the focal spot, which will displace the point P_i to $Q_i = \mathcal{T}(\bar{f}_i, F) \cdot P_i$. The aim is to choose an $\bar{F}$ which minimizes the net variation of $\Delta Q_i = Q_i - \mu_Q$. Through a series of computations, it can be shown that

$$\mu_Q = \mu_P + \frac{\mu_z}{\mu_{fz}}(\bar{F} - \bar{\mu}_f) \tag{6}$$

$$\Delta Q_i = (P_i - \mu_P) + [\frac{\Delta Z_i}{\mu_{fz}} - \frac{\mu_z \Delta f_{iz}}{\mu_{fz}^2}]\bar{F} + \frac{\mu_z \Delta f_{iz}}{\mu_{fz}^2}\bar{\mu}_f - \frac{\Delta Z_i}{\mu_{fz}}\bar{\mu}_f - \frac{\mu_z}{\mu_{fz}}\Delta \bar{f}_i \tag{7}$$

where $\mu_Q, \mu_P, \mu_z, \mu_{fz}, \bar{\mu}_f$ are the mean values of $Q_j, P_j, Z_j, f_{jz}, \bar{f}_j$; $Z_j = \mu_z + \Delta Z_j$ and likewise for $f_{jz}, \bar{f}_j$, where $j = 1 \ldots N$. In the above calculations, the second order terms either summed to 0 due to the independence of the variables or were too small in comparison. Our choice of $\bar{F}$ should be the one that minimizes the $variance(\Delta Q) = var(\Delta Q_x) + var(\Delta Q_y) + var(\Delta Q_z)$. It should be noted that F_z scales the whole space, *i.e.* a lower value will decrease the variance, implying that the choice of $F_z = 0$ forces $var(Q_z) = 0$ by forcing all $Q'_i s$ to lie on a plane. Thus $var(Q_z)$ does not provide sufficient constraints for F_z. We will first obtain F_x, F_y by minimizing the variance along x, y-axes (since there is no scaling in these directions), and then will compute F_z. Notice that the first term in Equation (7) is due to the relative movement in P, while the rest is due to an error in the calibration. Since we are interested only in the variance due to mis-calibration, we will ignore the variations in P. Minimizing $var(\Delta Q)$ and enforcing independence of f_{ix}, f_{iy} & f_{iz} gives

$$\bar{F} = \bar{\mu}_f - \frac{\Sigma_1^N \Delta f_{iz} \Delta \bar{f}_j}{\Sigma_1^N \Delta f_{iz}^2} \mu_{fz} = [\ \mu_{fx},\ \ \mu_{fy},\ \ 0\ ]^T \tag{8}$$

As expected, $F_z = 0$ from above. To compute F_z, we need to impose length preserving constraints. Thus if we measure a line segment of length l in each image, use Equation (2) to derive the net length error, the minimization implies

$$F_z = \mu_{fz}(1 - \frac{\Sigma_1^N \Delta f_{iz}^2}{N \mu_{fz}^2}) \ \sim \ \mu_{fz} \tag{9}$$

Thus $\bar{F} = \bar{\mu}_f$ (the mean), which is fairly intuitive and probably in common practice. Likewise, this particular choice of F_x, F_y is also a length preserving

constraint, *i.e.* it minimizes the error in lengths of line segments. Calibration error in ΔQ_i now reduces to $-\frac{\mu_z}{\mu_{fz}} \Delta \bar{f}_i$, which has a stable mean and low variance. Equation (10) gives a bound on the error when the assumed value of $\bar{F}$ is away from the mean $\bar{\mu}_f$ by a distance d. A 10 mm variation in the focal length ($var \sim$ 3 mm), $P's$ roughly at the iso-center having a depth variation of 100 mm and the assumed calibration unusually away from $\bar{\mu}_f$ by 50 mm still bounds the maximum error by 0.75 mm. Thus large and constant mis-calibration in many applications, might still provide sub-millimetric 3D quantitative measurements.

$$error \quad \leq \quad \frac{\sqrt{d^2 var(Z) + \mu_z^2 var(\|\bar{f}\|)}}{\mu_{fz}} \tag{10}$$

3   Phantom Experiments and Results

Validity of the Model: Equations (1) & (3) give the translation and rotation transformations as predicted by the affine model, the accuracy of which would furnish the validity of the model. We used the FTRAC fiducial (Figure 3), a small image-based fluoroscope tracking fiducial, which (given the calibration) can track a C-arm with an accuracy of 0.5 mm in translation and 0.65^o in rotation [7]. The fiducial was imaged using a Philips Integris V3000 fluoroscope and the true calibration read off the machine display. The images were not corrected for distortion. The pose of the fiducial (wrt to F_I) was first evaluated using the correct calibration, and then with the mis-calibrated parameters. The difference between the pose change predicted by the equations and the one computed using the non-linear pose estimation software, is displayed in Figure 2 (a) as a function of maximum calibration error. Even when mis-calibration is as high as 50 mm, the model can predict the rotation-axis with an accuracy of 4^o, amount of rotation under 1^o and translation under 1.5 mm. For extreme mis-calibrations the translation error linearly increases, while rotation is still stable. Thus the model seems to predict with an acceptable accuracy.

Accuracy of C-arm Tracking: The FTRAC fiducial was mounted on a 0.02^o accuracy rotational turntable, while the fluoroscope was kept stationary. The turntable was rotated by known precise amounts (ground truth) and images were taken. The relative poses were also computed using the pose estimation software. The accuracy in the estimation of C-arm motion is given by the difference between the computed relative pose and the true relative pose. The tracking accuracy is plotted in Figure 2 (b) as a function of mis-calibration. Even a high mis-calibration of 150 mm adds no *additional* error in C-arm motion estimation, fixing the value at 0.45 mm in translation and 0.6^o in rotation. An unusually high mis-calibration of 400 mm also only marginally decreases accuracy. Thus, mis-calibration does not increase the error of C-arm tracking .

3D Quantitative Reconstruction using Multiple Images: In addition to tracking a C-arm, it is equally important that multiple objects in the field of view (eg. vertebrae and screws) be reconstructed accurately relative to each
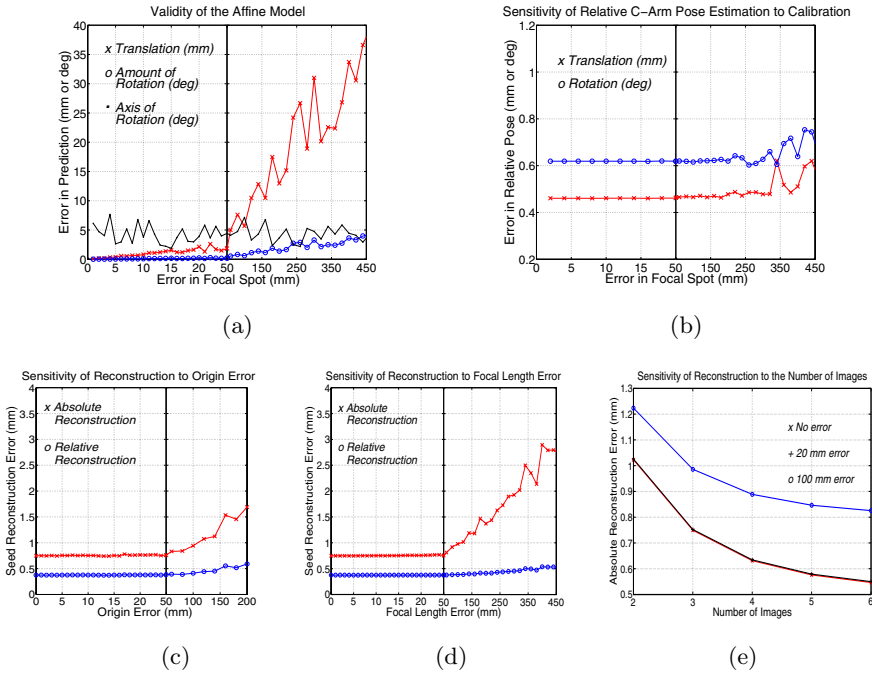
Fig. 2. Note the scale variation in x-axis. (a) An affine transformation is able to predict the movement of 3D objects due to mis-calibration; (b) C-arm tracking is insensitive to mis-calibration; 3D Reconstruction is insensitive to mis-calibration in (c) origin; (d) focal length up to 50 mm, beyond which it starts to linearly drift away from the tracking fiducial. Notice that the shape of the implant (relative err) is barely altered; (e) 3D reconstruction error decreases with an increase in images used.
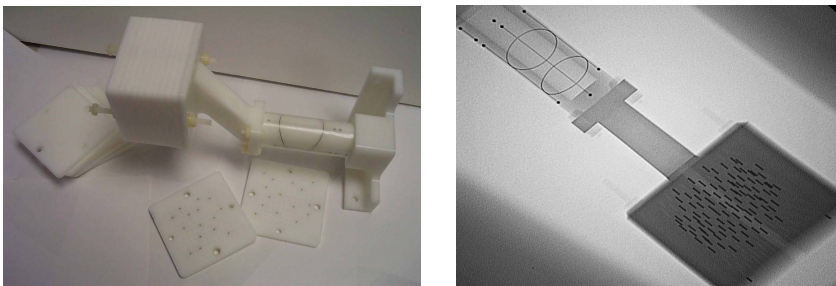
Fig. 3. An image of the seed phantom attached to the FTRAC fiducial (left). The seed phantom can replicate any implant configuration, using the twelve 5 mm slabs each with over a hundred holes. A typical X-ray image of the combination (right).

other. In order to validate our hypothesis that 3D reconstruction might not be sensitive to mis-calibration, we use an accurate acetol phantom (Figure 3) having 100 dummy radioactive seeds, approximating a brachytherapy implant

(Figure 3). The true 3D coordinate of each seed wrt the fiducial is known by rigid attachment. The C-arm is tracked using the FTRAC fiducial and the 3D seed coordinates are computed by triangulation (an algorithm called MARSHAL is used to establish correspondences). The difference between the computed and the true seed location gives us the 3D reconstruction error for each seed (wrt fiducial). The relative reconstruction error removes any consistent shift reflecting any change in shape. These errors are plotted as a function of mis-calibration in Figure 2 (c), (d). The reconstruction error is insensitive to mis-calibration in origin and focal length errors of up to 50 mm. The shape of the implant is stable even for large calibration errors. Figure 2 (e) shows a drop in reconstruction error as the number of images increase. Thus mis-calibration does not decrease reconstruction accuracy.

4   Conclusion

We modelled the the effects of mis-calibration on 3D reconstruction as an affine transform, and proved its validity experimentally. We have derived bounds on the amount of scaling, translation and rotation error. For pose dependant calibration, we proved that using the mean calibration minimizes the reconstruction variance. Phantom experiments with a radiographic fiducial indicate that C-arm tracking is insensitive to mis-calibrations. We also showed that mis-calibration up to 50 mm adds no additional error in 3D reconstruction of small objects, beyond which the reconstructed objects begin to drift wrt the fiducial, while still retaining the shape. In conclusion, a significant family of quantitative fluoroscopy applications involving localization of small markers can function without cumbersome on-line calibration. A constant loose calibration might suffice.

References

1. Hofstetter, R., Slomczykowski, M., Sati, M., Nolte, L.: Fluoroscopy as an imaging means for computer-assisted surgical navigation. CAS **4(2)** (1999) 65–76
2. Yao, J., Taylor, R.H., et al: A c-arm fluoroscopy-guided progressive cut refinement strategy using a surgical robot. Comput Aided Surg **5(6)** (2000) 373–90
3. Fahrig, R., et al: Three-dimensional computed tomographic reconstruction using a c-arm mounted XRII: correction of image distortion. Med Phys. (**24(7)**) 1097–106
4. Livyatan, H., Yaniv, Z., Joskowicz, L.: Robust automatic c-arm calibration for fluoroscopy-based navigation: A practical approach. In: MICCAI. (2002) 60–68
5. Kumar, R., Hanson, A.: Sensitivity of the pose refinement problem to accurate estimation of camera parameters. In: ICCV90. (1990) 365–369
6. Malis, E.: Visual servoing invariant to changes in camera intrinsic parameters. IEEE Transaction on Robotics and Automation **20** (2004) 72–81
7. Jain, et al: A robust fluoroscope tracking (FTRAC) fiducial. In: SPIE Medical Imaging; Visualization, Image-Guided Procedures & Display. (2005) 798–809

Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)

Ameet Kumar Jain[a)]
Department of Computer Science, Johns Hopkins University, Baltimore, Maryland 21218

Yu Zhou
Department of Mechanical Engineering, State University of New York at Stony Brook, Stony Brook, New York 11794

Tabish Mustufa
Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218

E. Clif Burdette
Acoustic MedSystems, Inc., 2110 Clearlake Boulevard, Champaign, Illinois 61822

Gregory S. Chirikjian
Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218

Gabor Fichtinger[b)]
Department of Computer Science, Mechanical Engineering, and Radiology, Johns Hopkins University, Baltimore, Maryland 21218

Intraoperative dosimetric quality assurance in prostate brachytherapy critically depends on discerning the three-dimensional (3D) locations of implanted seeds. The ability to reconstruct the implanted seeds intraoperatively will allow us to make immediate provisions for dosimetric deviations from the optimal implant plan. A method for seed reconstruction from segmented C-arm fluoroscopy images is proposed. The 3D coordinates of the implanted seeds can be calculated upon resolving the correspondence of seeds in multiple x-ray images. We formalize seed-matching as a combinatorial optimization problem, which has salient features: (a) extensively studied solutions by the computer science community; (b) proof for the nonexistence of any polynomial time exact algorithm; and (c) a practical pseudo-polynomial algorithm that mostly runs in $O(N^3)$ time using any number of images. We prove that two images are insufficient to correctly match the seeds, while a third image renders the matching problem to be of nonpolynomial complexity. We utilize the special structure of the problem and propose a pseudopolynomial time algorithm. Using three presegmented images, matching and reconstruction of brachytherapy seeds using the Hungarian algorithm achieved complete matching in simulation experiments; and 98.5% in phantom experiments. 3D reconstruction error for correctly matched seeds has a mean of 0.63 mm, and 0.9 mm for incorrectly matched seeds. The maximum seed reconstruction error in each implant was typically around 1.32 mm. Both on synthetic data and in phantom experiments, matching rate and reconstruction error achieved using presegmented images was found to be sufficient for prostate brachytherapy. The algorithm is extendable to deal with arbitrary number of images without any loss in speed or accuracy. The algorithm is sufficiently generic to provide a practical solution to any correspondence problem, across different imaging modalities and features. © *2005 American Association of Physicists in Medicine.* [DOI: 10.1118/1.2104087]

Key words: C-arm, fluoroscopy, seed matching, reconstruction, prostate brachytherapy, radiation planning

I. MOTIVATION AND BACKGROUND

With an approximate annual incidence of 220,000 new cases and 33,000 deaths prostate cancer continues to be the most common cancer in men in the United States.[1] For several decades, the definitive treatment for low risk prostate cancer was radical prostatectomy or external beam radiation therapy,[2] but low dose rate permanent seed brachytherapy (shortly brachytherapy thereafter in this document) today can achieve virtually equivalent outcomes.[3,4] The success of brachytherapy (i.e., maximizing its curative force while minimizing its co-morbidity) chiefly depends on our ability to tailor the therapeutic dose to the patient's individual anatomy. In contemporary practice, however, implant planning is based on idealistic preplanned seed patterns that, as 15 years of clinical practice has clearly demonstrated, are not achievable in the actual human body. According to a comprehensive review by the American Brachytherapy Society,[5] *the preplanned technique used for permanent prostate brachytherapy has limitations that may be overcome by intraoperative planning.* At the same time, continues the re-

port, *the major current limitation of intraoperative planning is the inability to localize the seeds in relation to the prostate*. There are excellent algorithmic and computational tools available today to optimize a brachytherapy treatment plan intraoperatively, thereby allowing for an improved dose coverage. These methods, however, critically require that the exact three-dimensional (3D) locations of the implanted seeds are precisely known with respect to the patient's anatomy.

Transrectal ultrasound (TRUS) imaging: Prostate brachytherapy is almost exclusively performed under TRUS guidance. While TRUS provides adequate imaging of the soft tissue anatomy, it does not allow for robust localization of the implanted brachytherapy seeds. Various researchers have tried to segment the seeds from TRUS images by linking seeds with spacers,[6] using x rays to initialize segmentation,[7] using vibro-accoustography[8] or transurethral ultrasound[9] as a new imaging modality, or segmenting them directly.[10] But even when meticulously hand-segmented, up to 25% of the seeds may remain hidden in ultrasound.[11] This necessitates the use of some other imaging method in intraoperative seed localization.

Fluoroscopy: The published history of C-arm fluoroscopy in brachytherapy originates[12] when it was first used as a solo guidance modality. Shortly after TRUS emerged as a primary image guidance modality, fluoroscopy became a secondary tool for gross visual observation. Mobile C-arms are ubiquitous in contemporary prostate brachytherapy, with approximately 60%[13] of the practitioners using it for qualitative implant analysis in the operating room. It is considered as the *gold standard* for intraoperative visualization of brachytherapy seeds. While several groups have published protocols and clinical outcomes favorably supporting C-arm fluoroscopy for intraoperative dosimetric analysis,[10,14–22] this technique is yet to become a standard of care across hospitals.

The ability to reconstruct and register the implanted seeds (that are visible in fluoroscope) to soft tissue anatomy (that is visible in TRUS) intraoperatively, would allow us to make immediate provisions for dosimetric deviations from the optimal implant plan. At the same time, quantitative use of fluoroscopy for dosimetric analysis has been hampered by a plethora of unresolved technical problems. The five major obstacles we face toward intraoperative dosimetry are: (a) C-arm distortion correction and calibration; (b) C-arm pose tracking; (c) seed segmentation; (d) seed matching and reconstruction; and (e) registration of C-arm to TRUS images.

Significant efforts have been made toward computational fluoroscopy guidance in general surgery,[23,24] developing various tools for distortion correction and calibration. However, C-arms available in most hospitals do not have encoded rotational joints, so one never knows where the fluoro shots are coming from relative to one another. We have addressed this issue by designing a fluoroscope tracking (hence-forth FTRAC) fiducial, which is a radiographic fiducial system creating a unique projection image from each direction.[25]

Various methods partially dealing with C-arm calibration in brachytherapy have also been proposed,[26–28] while some others have suggested that it is redundant.[29]

C-arm to TRUS registration: Attempts have been made to relate fluoroscopic images to soft tissue anatomy.[10,17,30–34] Nevertheless, further research is merited since they are susceptible to various kinds of errors. We address this issue by the use of the FTRAC fiducial. It is not only capable of tracking the C-arm, but also of registering the C-arm to TRUS by a predetermined placement.

Seed matching and reconstruction: We assume that the seeds are 3D points and that their image locations are known, i.e., we do not address automatic segmentation for which methods are available.[15,35–38] Three-dimensional coordinates of the implanted seeds can now be calculated from multiple x-ray images upon resolving the correspondence of seeds, which remains the focus of this paper. Formalization of the seed-matching problem results in a high complexity search space of the order 10^{150} and 10^{300}, from two and three fluoroscopic images, respectively. Hence previously proposed seed-matching approaches have predominantly been heuristic explorations of the search space, with no theoretical assurance on the accuracy of the answer.

The early attempts[39–41] toward seed matching used three *coplanar images* (coplanar images are those where the implant and the three x-ray sources are approximately in the same plane). The images were divided into variable width bands, formed by comparing coordinates along the rotation axis. Furthermore, in order to make the bands, it was assumed that the seeds are near the iso-center of the C-arm or at least have similar magnifications in all the images. These methods are prone to calibration errors and become ineffective as the number of seeds increases. These ideas were further extended by accommodating for patient motion,[42] and yet all the seeds could not be reliably reconstructed. Further geometrical constraints were imposed by assuming that some of the seeds are in a straight line[43] or on quadratic curves,[44] which due to seed migration seems to be an assumption not supported by compelling evidence.

The first step toward mathematical formalization came with the construction of a cost matrix,[45] where exhaustive matching gave the lowest cost solution. Though it eliminated extraneous assumptions, it required impractical computational resources. A greedy randomized algorithm,[46] tested with various cost-metrics, was suggested to reduce the run-time. This method gives a different output for each run and is typically iterated a few hundred times, choosing the sequentially lowest cost. Though this method might provide an answer close to the correct match, its randomization does make any claim on the number of iterations required or proximity of the final answer. Fast-CARS[47] is another variant, which significantly improved the computational complexity, where "*for each cost matrix an exhaustive match can be performed to obtain the best possible matches.*" It reduced the run-time from $O((N!)^2)$ to $O((A!)^2)$, where A is the average number of seeds in the band. Though it made the search faster, it still ran in exponential time. For example, if $A = 10$ then the num-

ber of computations would still be high at $O(10^{14})$.

Independently, a set of heuristic rules[16] were suggested for seed matching that attempted to reduce misclassifications. Simulated annealing[36] was proposed as an alternate technique to reach the global minimum. Another technique[28] was proposed that optimizes on seed positions and camera parameters, by generating simulated images and iterating them until they match the observed images. These optimization methods are prone to fall into local minima and may not be able to recover from them. A statistical simulation of seed reconstruction uncertainty was conducted,[48] but did not address the problem of seed matching. For completeness, we also mention that CT and MRI based techniques[49,50] are also proposed, but cannot be used intraoperatively and have poor resolution in the axial direction.

The matching problem is also prevalent in the computer vision community, where two-dimensional points are tracked and reconstructed to compute motion. Researchers have tried using noniterative greedy algorithms,[51] also incorporating spurious and hidden points.[52,53] Occlusion itself has also been a known problem.[54] These algorithms were optimized for a dense set of moving points, while specialized algorithms are used for sparse matchings,[55,56] which can also be used in pattern recognition[57] across images. These algorithms are usually catered to achieve real-time performance, as compared to a complete matching, and hence do not appear to be appropriate in a medical application.

The primary contribution of this work is a new theoretical framework for seed matching. The framework tackles issues of optimality and presents a practical algorithm that can be used. Moreover, the framework ensures a polynomial runtime of $O(N^3)$ on the algorithm, an improvement over previous methods. While it has been motivated by prostate brachytherapy, it is a general purpose correspondence algorithm that appears to be usable in many synergistic problems. Besides the aforementioned problem, we are exploring matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL) for the matching between 3D clouds of objects, fiducial-based registration in CT/MRI imaging, and calibration of tracked ultrasound. In these applications, MARSHAL appears to be a potent tool that can replace traditional iterative closest point and other gradient descent optimizations whose exit criteria are governed by custom-tuned thresholds. Section II describes the theoretical foundation and the proposed algorithm. Section III discusses the experiments and results. In Sec. IV, we discuss the shortcomings of the algorithm and the future work. Finally, we conclude in Sec. V.

II. MATERIALS AND METHODS

We convert seed-matching to network-flow-based combinatorial optimization. Our formulation has many salient features: (a) a global optimization using all seeds, as compared to a local seed clustering based approach; (b) exact solutions studied extensively by the computer science community; (c) addressing theoretically the achievable bounds by any algorithm; (d) guaranteed existence of a polynomial time solution
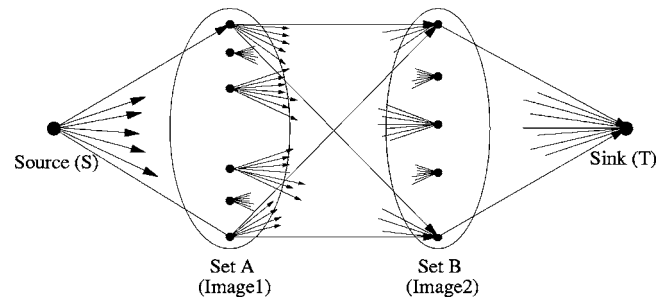
FIG. 1. The seed matching problem can be converted to a network flow graph. The best possible matching reduces to finding out the maximum flow with minimum cost.

for achieving the global minima for seed-matching from two images; (e) proof of the nonexistence of a polynomial time solution in case of more than two images; (f) a practical pseudo-polynomial algorithm using three images and mostly having an $O(N^3)$ claim on the space-time complexity; (g) an algorithm that does not become slower as the number of images is increased. A polynomial-time algorithm is one that runs in a number of computations, polynomial in the size of the input. Further details are available in the literature.[58] We do not assume any prior information (e.g., inserted seed positions), the value of which is questionable due to intraoperative seed migration and tissue deformation.

A. A generic network-flow-based formulation

A network flow formulation is created, where any flow in the network would represent a matching and the desired solution is the flow with minimum cost. Let N seeds be inserted and C-arm images I_1, I_2 be acquired. Let s_{ij} be the position of the ith seed in jth image. We construct a directed network as shown in Fig. 1.

Sets A and B, each with N nodes, represent the two images I_1 and I_2. While there are no edges within the set, directed edges (links) run from all vertices in set A to all vertices in set B. There are N links at source S, each link connecting to a node in A. Similarly each node in B is connected to sink T. The flow originates at S and ends at T, with each link allowing a flow of value 1 *or* 0, where 1 means that the edge is selected and 0 means that it is not. The problem is to efficiently compute a flow in the network that can achieve a total flow of value N.

It can be proved that any solution to the seed-matching problem is a solution to the flow problem and *vice-versa*. To have a net flow of N, each link connecting either the source or the sink has to support a flow 1. Now by the *conservation of flow* at each node, every node in set A will have to dispatch a unit flow to some node in set B. Moreover, each node in set B can accept only a unit flow, because any extra flow cannot be passed on to T and any deficiency would mean that T does not have a total flow of N units. The set of all links with nonzero flow provide a feasible matching. It can be verified that any matching of the seeds also provides a feasible flow. This proves that the flow problem is equivalent to the seed matching problem.

Simple combinations compute $N!$ feasible solutions to the seed-matching problem, giving rise to $N!$ feasible flows. To achieve the optimal solution, the link connecting seed s_{i1} to seed s_{j2} is assigned a cost $\mathcal{C}_{ij}$. The cost $\mathcal{C}_{ij}$ represents the likelihood of seed s_{i1} matching seed s_{j2}, with the cost being 0 if they match perfectly and ∞ (infinity) if they do not match at all. A popular example for the cost-metric is seed reconstruction accuracy. Further details are provided in Sec. II E. Any feasible flow has a net cost associated with it, the value of which is $\Sigma_{i=1}^{n}\Sigma_{j=i}^{n}\mathcal{C}_{ij}f_{ij}$, where f_{ij} is the flow in link ij and $\mathcal{C}_{ij}$ is the cost of sending a unit flow along that link. Thus the seed-matching problem is reduced to finding the flow with minimum cost, and can be written down as

$$\min \sum_{i=1}^{N}\sum_{j=1}^{N}\mathcal{C}_{ij}f_{ij},$$

$$where \quad f_{ij} \in \{0,1\}; \tag{1}$$

$$\sum_{j=1}^{N}f_{ij} = 1 \; \forall \; i \quad and \quad \sum_{j=1}^{N}f_{ij} = 1 \; \forall \; j.$$

The min-cost flow can be computed using cycle-canceling, successive shortest path, primal-dual, out-of-kilter or relaxation algorithms. These are all similar in their working and hence reducible to each other. Though they are straightforward to use, they run in pseudopolynomial time. The first weakly polynomial time algorithm was derived using an idea called scaling.[59] Capacity scaling, cost scaling, and double scaling algorithms that were developed on this idea are all weakly polynomial time algorithms. Repeated/enhanced capacity scaling and minimum mean-cost cycle algorithms achieve a strongly polynomial run-time. Alternately, linear-programming based techniques like simplex, interior point method or network simplex can also be used. Today, fast algorithms are available both freely[60] and commercially.[61] A comprehensive review of the above is available in the literature.[62]

We implemented the cycle canceling algorithm, which is the most straightforward of all the above. A feasible preflow is input as the current network flow. Then we generate the residual network, which is based on the difference between the network capacity and the current network flow, and search for negative cost cycles in the residual network. If there exists a negative cost cycle in the residual network, we decrease the cost by adding a flow along the negative cycle. Then we generate a new residual network based on this new flow in the network. We repeat this procedure until there is no negative cost cycle in the residual network. To find a negative cycle, we use the Bellman-Ford algorithm,[58] which runs in $O(Edges * Vertices)$ time. Thus the run-time for the cycle-canceling algorithm is $O(cost\;of\;the\;initial\;flow) \times O(Edges * Vertices)$, which for integer costs will be pseudopolynomial. One way to boost performance is to use the negative cycle with minimum weight instead of any negative cycle, but that computation itself is NP hard. So researchers have suggested choosing the minimum mean-
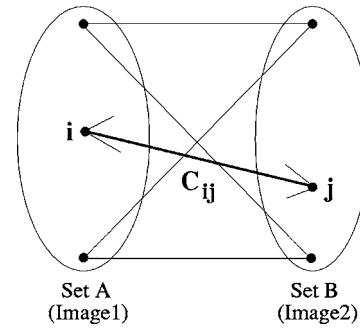

FIG. 2. Two image seed matching, modeled earlier as a min-cost max-flow problem, also reduces to the assignment problem. The best possible solution reduces to evaluating the minimum cost bipartite matching, which can be achieved in $O(N^3)$ time using the Hungarian algorithm.

cost cycle or canceling several node-disjoint cycles at once, both of which have strongly polynomial run-times.

Hence the seed-matching algorithm will have a polynomial run-time. Though strongly polynomial time algorithms are available, in practice the basic cycle-canceling algorithm itself was sufficiently fast. This is due to the favorable problem structure arising from a bipartite network and epipolar constraints of x-ray imaging. The worst case theoretical run-time above is for generic networks. It should be noted that a significantly faster $O(N^3)$ solution is possible if (a) all the nodes are available, (b) the number of nodes in both sets are equal, and (c) only unit flows are allowed. This is achieved by converting the generic network flow problem into a weighted bipartite matching problem, explained in the next section. Though we have implemented the cycle canceling algorithm, we have also tested and implemented a faster solution. However, this faster solution is not sufficiently generic to address the hidden seeds and other problems (Sec. IV).

B. Seed matching from two images

When all seed locations in the two images are known, the minimum-cost maximum-flow formulation reduces further to the specific problem of minimum-weight matching in bipartite graphs, also known as the assignment problem (illustrated in Fig. 2) and can be solved very quickly. The problem is to find a minimum weight subset of edges such that all the vertices are covered exactly once. A real-world problem could be the assignments of jobs to workers in a factory. Each worker (set A) can do some jobs (set B) at a certain cost. We seek an assignment of jobs, such that all jobs are completed, each worker gets exactly one job, and the total cost is minimized. The assignment problem is also formalized by Eq. (1).

The assignment problem is solved in $O(N^3)$ run-time by using the Hungarian algorithm.[62–64] Since the algorithm has been known for over four decades in the literature, we do not provide the full working of the algorithm, but instead just an outline. The $N \times N$ cost matrix $\mathcal{C}$ is constructed. The final aim is to choose exactly one element from each row (and column) such that the sum of the elements has the lowest attain-
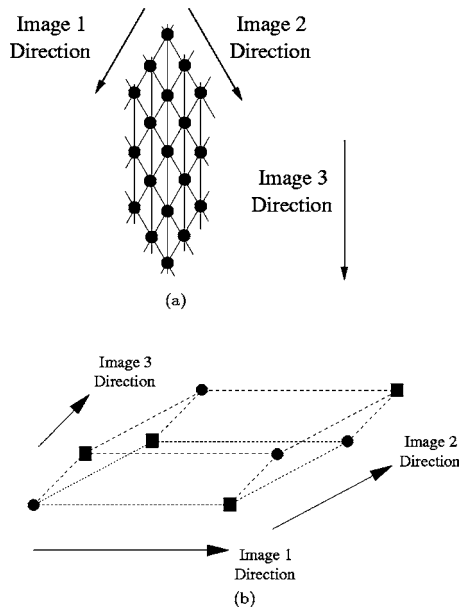
FIG. 3. (a) Three coplanar images also have seed constellation singularities. Multiple seed constellations will produce the same x-ray images. (b) Seed constellation singularity arising when three noncoplanar images are used. Two sets of reconstructions are viable for the same x-ray images. In fact this cube can be replicated to construct singularities that arise when more than three images are used.

FIG. 4. The three-image seed-matching problem reduces to the min-cost tripartite matching problem. Any see in a given set is connected to every seed in the other two sets. The problem is NP-hard and no polynomial time solution is possible to compute the min-cost matching.

able value. Thus, an equivalent matrix having at least one zero in each row (and column) is obtained by subsequent subtractions using the smallest element in each row (and column). This matrix is used to find a selection of zeros such that each row and column has exactly one zero. If this exists, then it provides the min-cost matching. If this does not exist, a *line covering* procedure is used to make an adjustment to the matrix and generate zeros in useful locations. This is iterated a maximum of $O(N)$ times (proof available in the literature) until a solution results. The locations of all zeros provide the minimum weight matching. Thus the Hungarian algorithm provides the matching with the lowest possible cost.

C. Theoretical foundations for uniqueness in matching

It is well known that due to singularities in projective imaging, reconstructing seeds from two images is inherently inaccurate. A third image is required to remove this singularity. Further investigations reveal singular constellations even with three images. Previously proposed algorithms use three coplanar images, i.e., where the three x-ray sources and the implant are all in the same plane. Figure 3(a) illustrates the arising singular constellations for this choice of images. Due to the small size of the prostate, the x rays entering it are nearly parallel, the average variation being around 2°. The symmetry establishes that there are multiple constellations producing the same x-ray images, with the number of alternate constellations increasing with the number of seeds. Errors are further amplified with inaccuracies in segmentation, C-arm calibration and tracking. Though we used symmetric
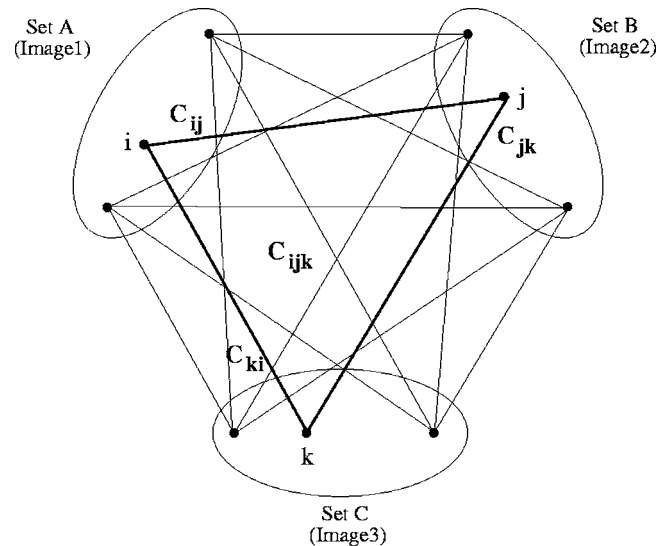
images, singularities can be shown in any three nonsymmetric coplanar images using the generic construction provided in the below.

Singularities though reduced, are not completely removed with noncoplanar images, as illustrated in Fig. 3. Since the x rays entering the prostate are nearly parallel, the three imaging directions define a parallelepiped, a 3D parallelogram. Four seeds sitting at the vertices of the parallelepiped can create a singularity. The likelihood of this happening increases for large implants, especially since the seeds are inserted in parallel straight lines. Thus a fourth image is advisable for large implants. Some singular constellations cannot be resolved by four images either. This constellation can be achieved by putting together two parallelepipes along the fourth imaging direction, which will result in two singular constellations with eight seeds each. This process can construct singularities with any number of images, if taken arbitrarily. *Theoretically* seven images can resolve all singularities for purely convex objects. Similar results can also be derived when other properties of the object are known.[65] In brachytherapy however, three to four noncoplanar images should *practically* be sufficient.

The above implies that a robust algorithmic framework for seed matching using three or more images is essential. A modification of our proposed framework reduces three-image matching to the tripartite matching problem,[64] as illustrated in Fig. 4 and Eq. (2) ($\mathcal{C}_{ijk}$ is the $N \times N \times N$ cost-matrix). Tripartite matching is similar to bipartite matching, except that it matches three sets instead of two. In addition, it assures a matching that is optimal, i.e., has a minimum cost. Multiple-image based seed matching similarly reduces to the multipartite matching problem. Though the bipartite problem is solvable in $O(N^3)$, the tripartite and multipartite problems are NP-complete. Moreover, finding the minimum weight tripartite matching is NP-hard. NP-complete are a class of dif-

ficult problems thought to be of nonpolynomial complexity. NP-hard problems are believed to be harder than NP-complete, where unlike NP-complete, even the validity of a given solution is not verifiable in polynomial time,[58]

$$\min \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{k=1}^{N} \mathcal{C}_{ijk} f_{ijk},$$

$$where \quad f_{ijk} \in \{0,1\}; \ \sum_{i=1}^{N} \sum_{j=1}^{N} f_{ijk} = 1 \ \forall \ k; \quad (2)$$

$$\sum_{j=1}^{N} \sum_{k=1}^{N} f_{ijk} = 1 \ \forall \ i \quad and \quad \sum_{k=1}^{N} \sum_{i=1}^{N} f_{ijk} = 1 \ \forall \ j$$

Using another reduction, three-image seed matching reduces to the min-weight *clique* problem[58] (a clique is a subgraph with edges between all the vertices). If the images are taken from arbitrary directions, then the orientation of epipolar lines is unconstrained. Since the implant has N densely packed seeds roughly shaped as an ellipsoid, any epipolar line connecting the projection of a seed to its x-ray source is expected to come in the vicinity of $O(N^{1/3})$ seeds in the 3D implant. Note that this estimate is an expected bound and could vary significantly across individual seeds. When two images are used, all seeds in the vicinity of the plane connecting the seed and the two x-ray source locations would become feasible. Thus, with unconstrained image orientations, each seed in an image *statistically* has an expected $O(N^{2/3})$ feasible correspondences in the other image. This implies that two-image matching has an expected $O(N^{5/3})$ feasible doublets. Similarly, it can be established that three-image matching has an expected $O(N^{4/3})$ feasible triplets.

This reduction has also been observed experimentally where two images resulted in about 2000 feasible doublets for a 100 seed phantom, while a third image reduced it to about 400 triplets. As the number of images is increased, the number of feasible tuplets will converge close to $O(N)$, implying a decrease in the number of alternate implants. This number can potentially reduce to exactly N, when many well-chosen images are used. It should be observed that even having only $(1+\epsilon) \times N$ feasible tuplets results in an exponential search space of $O(2^N \times (\epsilon-1)^N)$ for large ϵ and $O([2/(\epsilon-1)]^{[\epsilon-1]N})$ for small ϵ. A formal proof for the nonexistence of a polynomial-time algorithm can be shown by constructing a graph with all feasible tuplets as vertices, and compatibility between triplets as edges. The solution will be equivalent to finding a min-weight clique of size N. Finding a feasible clique in a general graph is itself NP-complete, while finding the clique with minimum cost[58] is NP-hard.

D. A practical algorithm for three-image seed matching

Though still an open problem in complexity theory, it is widely believed that no polynomial time algorithm can solve NP-complete or NP-hard problems.[58] The field of *approximation algorithms* could design quick algorithms that have

theoretical bounds on the error for both run-time and the final solution. Unfortunately, these algorithms typically work on generalized graphs, and may not necessarily incorporate the inherent structure of our problem. Hence we design our own algorithm that incorporates the physics of projective imaging. We propose a practical solution for matching and reconstruction of brachytherapy seeds using the Hungarian algorithm, abbreviated henceforth in the paper as MARSHAL.

The intuition behind MARSHAL can be summarized into: (a) the original tripartite matching can be projected into *inspired* bipartite matchings; (b) s_{i1} matches s_{j2} only if they have a counter-part in the third image; (c) though a low $\mathcal{C}_{ij}$ did not force $\mathcal{C}_{ijk}$ to be low for some k, a *low* $\mathcal{C}_{ijk}$ does force $\mathcal{C}_{ij}$, $\mathcal{C}_{jk}$, and $\mathcal{C}_{ki}$ to be *low*. It is very common in the two-image case for the epipolar line of s_{i1} to pass through s_{j2}, even when s_{i1} and s_{j2} are severely mismatched. Thus in a matching between images 1 and 2, $\mathcal{C}_{ij}$ should be low if and only if there exists a k such that $\mathcal{C}_{ijk}$ is low. If $\mathcal{C}_{ijk}$ if high for all k, then $\mathcal{C}_{ij}$ should be given a high value. This assignment of $\mathcal{C}_{ij}$ is different from the purely two-image case, where $\mathcal{C}_{ij}$ could be low even for a strong mismatch. Thus $\mathcal{C}_{ij}$ needs to incorporate information from the third image to remove inherent two-image singularities. Since the solution of bipartite matching should as close as possible to that of the tripartite matching, the $L_{-\infty}$ projection (minimum/best value) is used.

$$\mathcal{C}_{ij} = \min\{\mathcal{C}_{ijk}/k = 1, \dots, N\}. \quad (3)$$

MARSHAL *projects* the original tripartite problem into three distinct bipartite problems by the appropriate projection of the costs. This is similar to projecting the *minimum value* along the rows/columns of the three-dimensional cost matrix $\mathcal{C}_{ijk}$ to obtain three two-dimensional cost matrices. The theoretical framework of bipartite matchings allows the computation of the best possible matching for each bipartite problem in $O(N^3)$ run-time. The solutions of the bipartite matchings are then *integrated* to obtain a solution to the original tripartite problem. Hence MARSHAL acts as a bridge between a purely theoretical framework and an inherent structure in the problem.

The flowchart for MARSHAL is shown in Fig. 5. Using Eq. (3), the respective cost matrices for the three image pairs are obtained. These are used to achieve three independent bipartite matchings $M_1\langle i,j,-\rangle$, $M_2\langle -,j,k\rangle$, and $M_3\langle i,-,k\rangle$. Loops are created next, i.e., if $\langle i1,j1,-\rangle,\langle -,j1,k1\rangle$, $\langle i2,-,k1\rangle,\langle i2,j2,-\rangle\langle i1,-,km\rangle$ are matchings, then $\langle i_1,j_1,k_1,i_2,j_2,k_2,\dots,k_m,i_1\rangle$ is a loop of size m, each loop ending on the seed it started from. For example, if for the first seed the matchings are $M_1\langle 1,1,-\rangle$, $M_2\langle -,1,1\rangle$, $M_3\langle 1,-,1\rangle$, then $\langle 1,1,1,1\rangle$ is a loop of size 1 representing *seed1*. The majority of the loops will have a size of 1, which can be declared matched. This claim is also experimentally validated in Sec. III C.

Sometimes matchings get flipped, resulting in loops of size greater than 1. For example, if s_{13} had instead matched s_{21} in the above-noted example. i.e., if the matchings
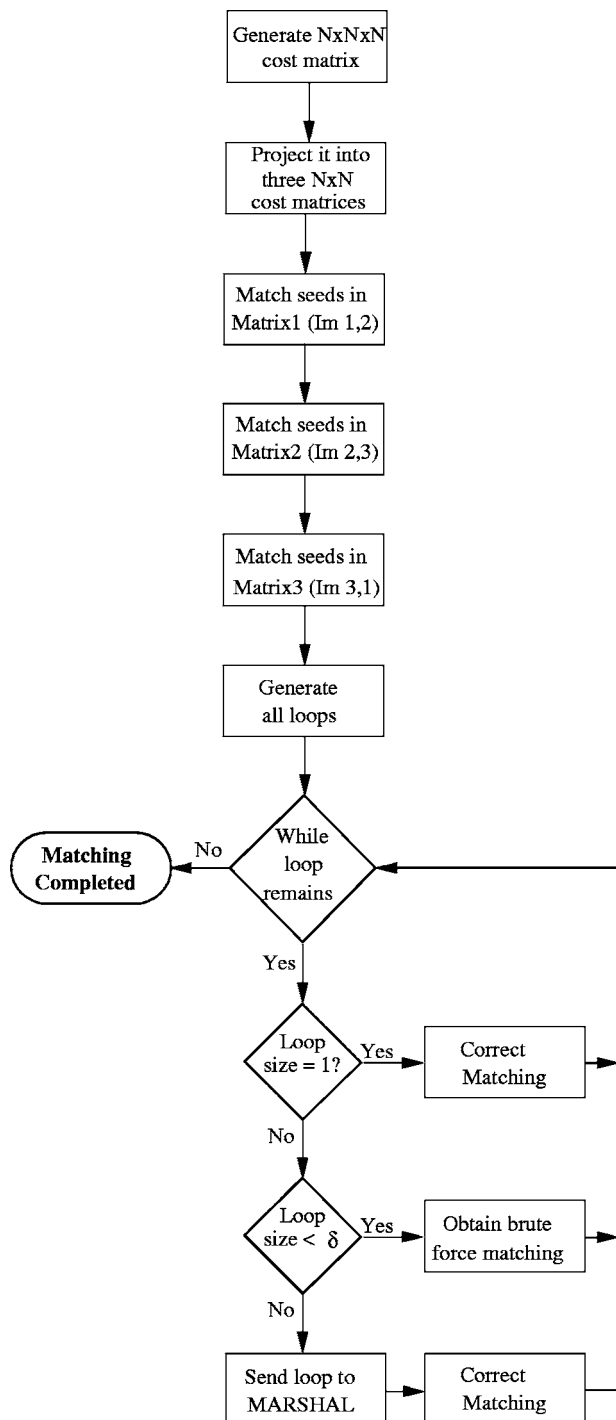
FIG. 5. A flowchart explaining the working of MARSHAL. It runs in $O(N^3)$ time for *good data sets*.

were $M_1\langle 1,1,-\rangle$, $M_2\langle -,1,1\rangle$, $M_3\langle 2,-,1\rangle$, $M_1\langle 2,2,-\rangle$, $M_2\langle -,2,2\rangle$, $M_3\langle 1,-,2\rangle$, then $\langle 1,1,1,2,2,2,1\rangle$ would be a loop of size 2 representing *seed*1 and 2. Since each loop *by definition* is a self-contained subset of seeds with no relation to other loops, they only need to be matched among themselves to obtain the correct final answer. These are typically small loops, the optimal match for which can easily be obtained by a brute-force search with $O(m! \times m!)$ run-time.

Theoretically, for error-prone data the size of the loops could grow large, though practically we never saw it happen. Nevertheless, brute force search is done only for loops of size $m < \delta$, which is a predetermined threshold based on the largest m that achieves a very fast (<1 s) run-time. Loops of size greater than δ are *recursively* broken down using MARSHAL, which would run three bipartite matchings within this subset and return the correct match. In effect, all *large* loops will be broken down into smaller and smaller loops, until everything is matched. Thus MARSHAL achieves the correct seed matching and still practically runs in $O(N^3)$, a significant improvement when compared to previous methods.

It should be noted that there can be another way to compute loops in the above algorithm. If $\langle i,j,-\rangle$ match after the first bipartite matching, then their counterpart k in *Image*$_3$ could also be chosen. We do not do this, since the first bipartite matching does not ensure that all $k's$ are chosen, to solve which the full tripartite problem would have to be solved. Moreover, $\langle i_1,j_1,-\rangle$ and $\langle i_2,j_2,-\rangle$ can map on the same k, or some k can be left untouched after the first bipartite matching. Thus the first bipartite matching attempts to compute as accurate as possible a correspondence in the first two images, while the next two matchings will later solve for any inconsistencies.

Note that the computation in Eq. (3) takes $O(N^3)$ time. We modify it to run in $O(N^2)$ time. To compute $\mathcal{C}_{ij}$, s_{i1} and s_{j2} are used to reconstruct the 3D point S'_{ij}, which is then projected on *Image*$_3$. Let k be the closest point in *Image*$_3$ to this projection. Now S_{ijk} is computed and its projection error (PE) calculated, which is assigned to $\mathcal{C}_{ij}$. Though this $\mathcal{C}_{ij}$ does not strongly satisfy Eq. (3), it weakly satisfies it. Any consistent $\langle i,j\rangle$ will choose the correct k, while an inconsistent $\langle i,j\rangle$ does not have any correct counter-part k, and will anyway be eliminated in the bipartite matching. Thus our $\mathcal{C}_{ij}$ exactly satisfies Eq. (3) in relevant cases, and will become irrelevant otherwise.

E. Cost-metric

A comparison proves that some metrics are more effective than others.[46] The metric should incorporate all the available information, making the global minimum sharper and the algorithm robust. Hence $\mathcal{C}_{ijk}$ should ideally represent the *inconsistency* among seeds s_{i1}, s_{j2}, and s_{k3}, with 0 indicating least inconsistent. The basis for most popular choices of $\mathcal{C}_{ijk}$ are *reconstruction accuracy* (RA) and *projection error* (PE). We do not formally compare the various metrics that we have tried, but instead provide the intuition behind the preferred one.

RA, used by most researchers, is computed by first calculating the equation of the three lines that join each projection to its respective x-ray source. Due to various errors these lines never intersect, creating a need for a symbolic intersection. As explained in Sec. II F, a closed form solution for RA that minimizes the L_2 norm of the distance vector is used in this publication. Though they do not have an easy closed form solution, L_1, L_3, L_∞, or other norms can also

be minimized ($\|v\|_{L_i} = \sqrt[i]{\Sigma_k v_k^i}$). The L_∞ norm would the most robust, since the higher the order of the norm, the more sensitive it is to the smallest of inconsistencies. Thus, it measures the accuracy with which the 3D point can be reconstructed. PE, an alternate metric, can be computed by extending RA. The 3D reconstructed point can be projected back into each image, and the mean distance between the projected location and the observed location of the seed constitute PE. The symbolic intersection of the 3D point can be computed by minimizing any choice of norm ($L_1 - L_\infty$), and could possibly affect a change in the performance. We have minimized the L_2 norm to compute the 3D intersection, and used its PE as the metric.

In our experience, PE fares significantly superior as a cost-metric than RA. Though PE for each seed is its RA magnified, different seeds will have different magnifications (depending on depth). If two seeds have similar RA but differing PE, the seed with lower PE should ideally be chosen by the optimization. On the other hand, an algorithm minimizing RA as a metric will not necessarily choose a correspondence that minimizes the deviation of reconstruction from observation. Thus PE is a stronger reflection of the observed data and leads to better convergence. In general, any cost-metric that directly measures the deviation from the observation performs superior to a metric that does not. A significant improvement in matching was observed by switching from RA to PE. A combination of the above-mentioned ideas is to choose the intersection so as to minimize the L_∞ norm of PE. Thus the most sensitive norm (L_∞) is minimized on the observations (seed locations on image), to compute the 3D intersection. Unfortunately, it requires a nonlinear optimization due to a lack of a closed form solution. Though we have not yet used it, this metric is expected to give the best performance.

F. Seed reconstruction

To compute $\mathcal{C}$, we need to compute the 3D intersection of the corresponding straight lines in space. Due to various errors these straight lines never intersect, forcing us to compute a symbolic 3D intersection point. The symbolic intersection is typically defined as the global minimum of an error function. Here we propose a simple and quick method that minimizes the L_2 norm of Euclidian distance from the intersection point to the lines.

Let the total number of 3D straight lines be m, with *line i* defined as having unit direction l_i (a_i, b_i, c_i) and a point p_i on it, as shown in Fig. 6. Let P (x, y, z) be the representative intersection of these m lines. Let d_i be the Euclidian distance of P from *line i*. Thus by definition, P achieves the minimum L_2 norm for the vector $(d_1, d_2, \ldots, d_m)$. In other words, we need to find a P such that it minimizes a function $\mathcal{F}$. Now, it can be easily computed that
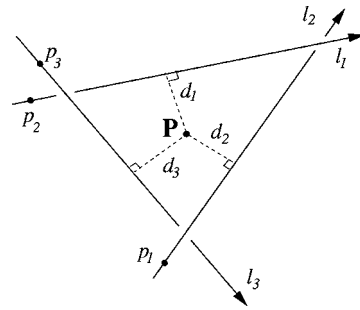

FIG. 6. Three matching points, in general, do not intersect due to various errors in segmentation, C-arm pose, and calibration. A symbolic intersection is calculated by finding the point with minimum sum of square distances from the lines.

$$\mathcal{F} = m \times \mathrm{RA}^2$$

$$= \sum_{i=1}^{m} d_i^2$$

$$= \sum_{i=1}^{m} \|(P - p_i) \times l_i\|^2$$

$$= \sum_{i=1}^{m} (P - p_i)^T \begin{bmatrix} b_i^2 + c_i^2 & -a_i b_i & -a_i c_i \\ -a_i b_i & a_i^2 + c_i^2 & -b_i c_i \\ -a_i c_i & -b_i c_i & a_i^2 + b_i^2 \end{bmatrix} (P - p_i)$$

$$= \sum_{i=1}^{m} (P - p_i)^T A_i (P - p_i), \tag{4}$$

where T represents the transpose. Now P minimizes $\mathcal{F}$. Hence we can write, $\partial \mathcal{F} / \partial x = \partial \mathcal{F} / \partial y = \partial \mathcal{F} / \partial z = 0$, which leads us to

$$P = \left[\sum_{i=1}^{m} A_i \right]^{-1} \times \left[\sum_{i=1}^{m} A_i p_i \right]. \tag{5}$$

This is the final seed coordinate. We can see that it can be computed very quickly by a few summations followed by a 3×3 matrix inversion. It should be noted that P is chosen so as to minimize RA.

G. Seed matching from four or more images

In some rare cases of extremely large/dense implants, an additional fourth image may be desired. One approach is to match all possible image pairs and perform a combinatorial search on the loops that are formed. Though this is accurate, the run-time complexity is dependent on the number of images, making the algorithm slow. It was earlier observed that the number of feasible m-tuples rapidly comes closer to N. This property can be used to extend MARSHAL. The algorithm chooses three images at random as primary and the rest ($m-3$) as secondary. Using the primary images, it runs exactly as described in Fig. 5, except that each bipartite problem is a projection of the m-partite problem. In other words C_{ij} incorporates information from all the m images, as shown as follows:
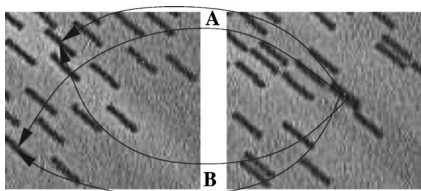
FIG. 7. Two seeds nearly overlap in the right image. Though both matchings A and B reconstruct the seeds in the correct location, only one of them is theoretically correct. We define the other as mismatched, so as to make the seed matching error analysis extremely rigorous.

$$\mathcal{C}_{ij} = \min\{\mathcal{C}_{i_1 i_2 \ldots i_m}/i_1 = i; i_2 = j \text{ and } i_3, \ldots, i_m \in \{1, 2, \ldots, N\}\}.$$

$$(6)$$

Extended-MARSHAL will still run in $O(N^3)$ time. In comparison to three-image matching, using four images will in fact have a faster run-time due to a decrease in loop-size and due to a better conditioned bipartite matching. An improvement in accuracy is also registered owing to the inherent structure in the additional information.

H. Guidelines for performance evaluation

Since MARSHAL is essentially a correspondence algorithm, the % of correctly matched seeds (mean, STD) is the optimal metric to evaluate its performance. However, it should be noted that due to the *overlap-problem*, the % of declared mismatches would always be an upper-bound. As illustrated in Fig. 7, both matchings **A** and **B** give the correct 3D seed coordinates in a *practical sense*, though only matching **A** is correct in a *theoretical sense*. Owing to our rigid error analysis criterion, we declare **B** to be an incorrect matching, in spite of it being correctly matched in a practical viewpoint. This will be evident from the low reconstruction errors for mismatched seeds, which should be comparable to that of correctly matched seeds.

Error analysis is sometimes also done using PE or RA. As theoretically described in Sec. II C, a mismatch might still result in a low PE/RA. This makes performance evaluation based on PE/RA inaccurate. If available, reconstruction error (RE), the 3D distance between the true and the reconstructed seed location, is an excellent tool for measuring performance. To measure any systematic bias in implant reconstruction or any change in shape, we use the relative reconstruction error (relative RE). It is the 3D error when the reconstructed implant is compared to a shifted (translated, rotated) version of the ground truth implant.

Typically RE is averaged for all the seeds, leading to potential information misrepresentation. For example, a 100-seed implant can have 90 seeds correctly matched with 0.25 mm RE, and 10 severely mismatched seeds with 5 mm RE. Though the *average* RE in this case is only about 0.7 mm, the 10 mismatched seeds could be hazardous. Thus a small average RE alone does not necessarily imply a good matching. *A separation of matched and mismatched seeds is necessary.* Note that the distinction between matched and mismatched seeds should never be established by using PE/

RA, but only by an accurate ground truth. RE/PE/RA in addition to the above, also reflect the accuracy of C-arm tracking and seed segmentation. In this light, MARSHAL is primarily evaluated using % matching and RE for mismatched seeds.

Beside establishing tools for performance evaluation, it is also important to find factors that could adversely affect a seed matching algorithm. This is crucial, since no polynomial algorithm can assure optimality for more than two images. The primary factors are the implant seed density and the size of the prostate. The higher the density or bigger the prostate, the greater the possibility of mismatches. Besides, most available C-arms are of moderate quality, with little reliability of the system supplied parameters. This includes encoder readings (pose) and parameters like image warp, focal length, and image origin. Another factor that could affect performance is the angle of separation between the images. Thus sensitivity to these factors should be considered when evaluating any new algorithm.

III. RESULTS AND DISCUSSION

A. Two-image versus three-image matching

By converting the two-image matching problem to a bipartite matching problem, we have assured that the minimum-cost match from all the possible $N!$ matches can be computed in $O(N^3)$ run-time. This proves that there cannot exist a better match, implying that no other algorithm can produce a match with a lower cost. It should be noted that the best possible match is not representative of the correct match, which is due to the fact that two images are intrinsically insufficient and do not contain the requisite information. Furthermore, the minimum-cost match is dependent on the choice of cost-metric, i.e., a superior choice for the cost-metric will result in a superior match. Hence we have proved that given two images and a choice for the cost-metric, MARSHAL computes the best possible match (global optima) in polynomial time. Furthermore, we have shown that information from two images is inherently ill-posed and that the best match does not correspond to the correct match.

This is also validated using quantitative experiments. Synthetic C-arm images of clinically realistic brachytherapy implant scenarios were created, and the reconstructed and actual seed locations were compared. Implants with the number of seeds varying from 60 to 150 in a 50 cc prostate were tested. Based on our experience, the focal length was chosen to be 1000 mm, the pixel size was chosen to be 0.25 mm and the center of the image was chosen to coincide with the imaging origin. The region of use was at about a distance of 2/3rd the focal length from the x-ray source. The C-arm images were all contained inside a 15° cone around the AP axis. To reflect the nature of the best possible results, no noise was added in these simulations. Table I shows the performance on this synthetic data.

It can be noticed that two images are clearly insufficient, correctly matching only 85.2% of the seeds even in the absence of any noise. The inherent singularity in the two-image case is indicated by the low PE and high RE for the mis-

TABLE I. MARSHAL performance using synthetic data. Two images are clearly insufficient, which is clear from the low projection error yet high reconstruction error for the mismatched seeds. All seeds are matched when three images are used.

		Number of seeds									
		Two images					Three images				
		60	80	100	120	150	60	80	100	120	150
Matching rate (%)	Match	91.3	90.1	80.9	82.8	80.8	100	100	100	100	100
Reconstruction error (mm)		0.18	0.15	0.18	0.20	0.21	0.07	0.07	0.07	0.07	0.07
Reconstruction error (mm)	Mismatch	22.1	21.4	22.1	23.6	21.9	⋯	⋯	⋯	⋯	⋯
Projection error (mm)		0.03	0.04	0.04	0.04	0.04	⋯	⋯	⋯	⋯	⋯

matched seeds. Thus, when only two images are used, there exist alternate implants that produce the same x-ray images (low PE of 0.04 mm), but reconstruct the mismatched seeds at the wrong 3D location (high RE of 22.2 mm). Another way to look at the results is that the desired matching of the seeds has a cost associated with itself. This cost is close to 0 (in practice it is always a little more than 0 due to pixelization errors). Due to inherent singularities, a significant number of alternate implant constellations also have a cost very close to 0. Theoretically, all these are viable implants. Nevertheless, even a pixel size of 0.25 mm results in many of these alternate implants having a cost lower than the desired implant, which is then chosen by the algorithm. Thus, any algorithm which is minimizing the error from observed data will choose an incorrect matching. This is not the fault of the matching but rather a deficiency of the input. The suggested approach, therefore, would be to increase the number of images. Similar conclusions for two-image matching can also be drawn from phantom experiments, complete results of which are available in Sec. III C.

In comparison, using a third image matched all seeds correctly on noise-free data. Thus while two images are insufficient, three are practically adequate. As mentioned in Sec. III C, three-image matching reduces to the tripartite matching problem. Multiple-image based seed matching similarly reduces to the multipartite matching problem. Though the bipartite problem was solvable in $O(N^3)$, the min-cost tripartite and multipartite problems are NP-hard, meaning no polynomial algorithm will be able to prove that its solution is the best achievable. Thus this proves that there cannot exist any polynomial time algorithm which can solve the matching problem using three or more images. It should be noted that algorithms providing practically acceptable solutions (without any claim on absolute optimality) can still exist.

B. Simulations

Studies were conducted on synthetic images to analyze the effects of various governing parameters. MATLAB software was created to model x-ray imaging. Given C-arm parameters and implant details, it generated synthetic images

and exact locations of seeds in the images. Random error was modeled using a uniform probability density function, i.e., a 1 mm error means that a maximum error of *magnitude* 1 mm was added, having a uniform probability distribution. The C-arm geometry was the same as in the previous section. Individual parameters were changed to understand the sensitivity of MARSHAL. It should be noted that only three images were used to evaluate the correspondences in all the simulation experiments. Using four or more images is expected to improve performance, though not evaluated in simulation experiments.

Since MARSHAL is essentially a correspondence algorithm, to evaluate performance on simulated data, we computed only the % of correctly matched seeds (mean, STD). RE for matched seeds is not plotted, as it is a function of the particular methods used for C-arm tracking and segmentation, which are not the focus of this paper.

Segmentation: Seed segmenting from the x-ray images is one of the most important sources of error. To study the sensitivity of MARSHAL to segmentation, segmentation error ranging from 0 to 2 mm in steps of 0.25 mm was added to each seed coordinate in the synthetic data sets. The computed correspondence was compared to the correct known correspondence, and the % of correctly matched seeds was evaluated. The seed density was varied from 1.5 to 1.9 seeds/cc in steps of 0.1, while the prostate volume was varied from 45 to 55 cc in steps of 5 cc. The averaged results (mean and STD) from a total of 54 000 iterations are plotted in Fig. 8(a).

The results indicate that even with segmentation errors as high as 1 mm, the mean and STD for % matching is better than 97% and 2%, respectively. With 1–2 pixel segmentation error, the average matching rate is practically 100% in all data sets. Segmentation error needs to increase beyond 1.5 mm for the mean matching rate to fall below the 95% mark. Moreover, it should be noted that the declared mismatches are predominantly due to the overlap problem, which becomes clear from the low RE for mismatched seeds. It has an average of 0.52 mm for all the data sets. The maximum is attained for data sets with 2 mm segmentation error
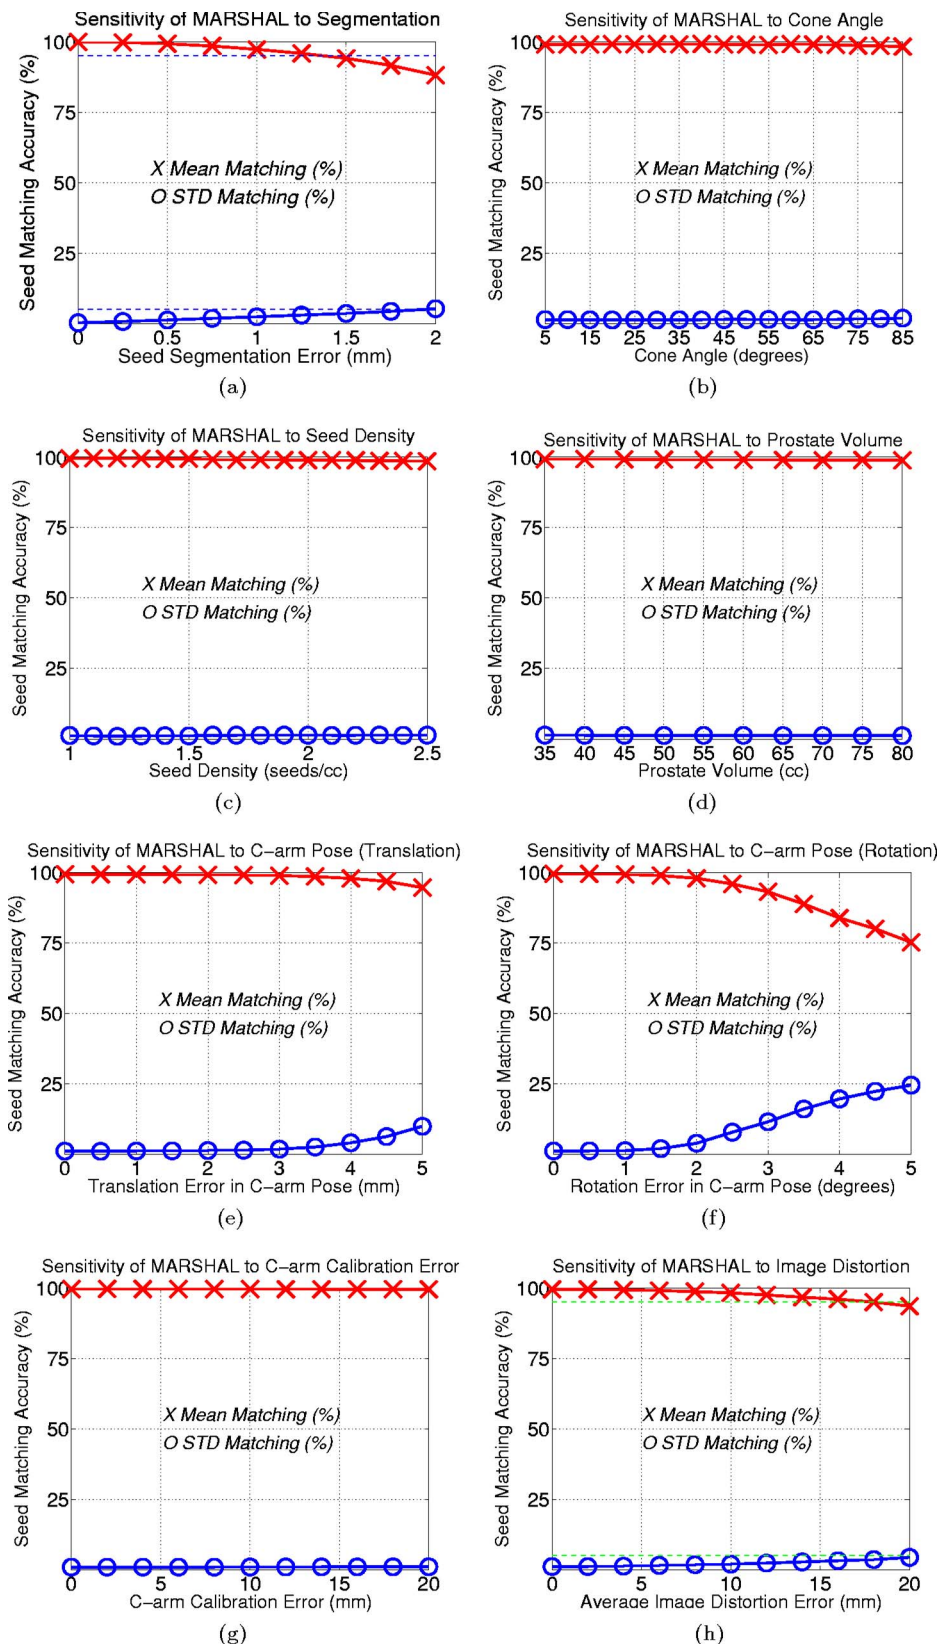
FIG. 8. Error in the performance of MARSHAL: (a) segmentation error; (b) cone angle variation; (c) seed density variation; (d) prostate volume variation; (e) C-arm pose translation error; (f) C-arm pose rotation error; (g) C-arm calibration error; and (h) C-arm distortion error.

at 1.51 mm, which is natural due to the presence of noise in segmentation, as compared to an error in matching. MARSHAL requires seed segmentation errors to be under 3–4 pixels (~1 mm) for optimal performance. For the re-

mainder of the simulations, we fixed a segmentation error of 0.5 mm (1–2 pixels) to be added to all the datasets.

Seed density and prostate volume: Depending on the radio-activity of the chosen seed, the seed density is com-

puted. From some preliminary calculations, seed density typically ranges from 1.5 to 2 seeds/cc.[66] An increase in seed density could make the matching process more difficult, since the seed projections in the x-ray image gets more and more congested. Thus it is extremely necessary to evaluate the sensitivity of any algorithm to seed density. In a similar manner, a larger prostate would need a greater amount of dose, resulting in a greater number of seeds. An increase in the number of seeds can potentially be of concern as the number of feasible matches would increase. Matching algorithms are expected to be less sensitive to prostate volume in comparison to seed density. Nevertheless, they should be evaluated in view of volume changes.

The seed density was changed from 1 to 2.8 seeds/cc in steps of 0.1 and the prostate volume was changed from 35 to 80 cc in steps of 5. As mentioned before, segmentation error was kept constant at 0.5 mm. A total of 23 750 data sets were evaluated. Averaged results for matching rate, as a function of seed density and prostate volume are plotted in Figs. 8(c) and 8(d), respectively. It can be noticed that both the mean and STD are stable throughout with a variation of less than 1%. The mean matching rate is 99% or above, with a few seeds in some cases being counted as mismatched due to the overlap problem. Thus MARSHAL can be declared robust to variations in seed density and prostate volume.

Separation angle: The three or more x-ray images of the implant are typically taken with a rotation motion of the C-arm. Some earlier work had made certain assumptions on the path of the C-arm, while some other work used orthogonal images to do the reconstruction. It is desirable to not have any constraints on how the images are taken. Moreover, in a clinical setting the C-arm would have only a limited mobility inside a 25° cone around the AP axis. Thus, it is necessary to validate an algorithm to the angular separation between the images. We tested the algorithm on various data sets with images taken on a cone around the AP axis. The implant was kept close to the center of the cone, which was in proximity of the iso-center. The cone angle was varied from 5° to 85° in steps of 5°. Six images were taken on each cone and all 20 three-image combinations used. Seed density was varied from 1.8 to 2 seeds/cc, prostate volume from 45 to 55 cc, and segmentation error was constant at 0.5 mm.

Averaged results from a total of 77 400 data sets are plotted in Fig. 8(b). Mean matching rate is fairly constant remaining always over 99%. There does not seem to be any variation in % matching due to a variation in the angular separation of the images. Thus it can be concluded that MARSHAL is robust to image-separation, and any three images in the available workspace should suffice. It should be noted however that, though MARSHAL can match the seeds correctly it does not have a bearing on RE, which for each seed will be a function of cone angle, C-arm tracking and seed segmentation. Thus a larger image separation is desired since it improves RE.

C-arm pose tracking: The accuracy of C-arm tracking has a definite effect on the ability to correctly match the seeds. We separately evaluate the performance of MARSHAL to errors in C-arm translation and rotation. Statisti-

cally uncorrelated pose errors were added to all three images. Errors in translation varied from 0 to 5 mm in steps of 0.5 mm, while in rotation varied from 0° to 5° in steps of 0.5°. It was observed in our fluoroscope tracking (FTRAC) fiducial that translation errors in depth are always significantly greater than those parallel to the plane[25] (by a factor of at least 5). This statistical bias in translation was incorporated in the generation of the datasets. There was no significant bias observed in the rotation errors.

Averaged results for translation and rotation from a total 44 000 data sets are shown in Figs. 8(e) and 8(f), respectively. It can be noted that translation errors up to 3 mm can still correctly match over 99% of the seeds. This quickly drops to 95% when the error reaches 5 mm. When compared to translation, MARSHAL is more sensitive to errors in rotation. Errors up to 1.5° in rotation correctly match over 99% on the seeds, with the performance dropping quickly when the error is greater than 3°. It can be concluded that MARSHAL is accurate and robust when the C-arm is tracked to an accuracy of 3 mm in translation and 2° in rotation. This is significantly larger than accuracies offered by the FTRAC (0.56 mm in translation and 0.33° in rotation)[25] or other fiducials.[34]

C-arm calibration: Quantitative fluoroscopic reconstruction requires accurate calibration of the C-arm imaging model parameters. Since the parameters vary from pose to pose, this is typically done for every pose from which the image is taken. The five intrinsic camera parameters are the pixel sizes (two parameters) and the 3D location of the x-ray source with respect to the image (three parameters). Since pixel sizes are fixed throughout the life of the C-arm, we evaluate MARSHAL only with respect to the focal spot. Independent and uncorrelated calibration errors of up to 20 mm were added to all three images, and the averaged % matching from 102 500 data sets are plotted in Fig. 8. It can be seen that the matching rate stays practically at about 100%. Thus MARSHAL is robust to C-arm calibration. These results triggered us to question the relevance of C-arm calibration in brachytherapy, leading to both theoretical and experimental bounds on the error.[29]

C-arm image distortion: Most C-arms exhibit a significant amount of nonlinear distortion in the images. Moreover, the amounts of distortion vary with pose, time, and location. Distortion being a global and nonlinear transform, consistently shifts the seed coordinates in the image, adding an error to any algorithm that relies on epipolar constraints. Thus distortion correction needs to be accommodated for in every image, which is a cumbersome process. If any matching algorithm is stable in the presence of distortion, then a one-time distortion correction procedure might also suffice. MARSHAL solves the matching problem globally, as compared to decreasing the complexity to smaller epipolar constraint based subsets, and is well suited to work on distorted data sets.

MARSHAL was tested for sensitivity to distortion. A randomly generated fifth degree radial distortion function was used to distort each image, and the matching computed on this data set. The implants were not confined to be in the

(a)   (b)
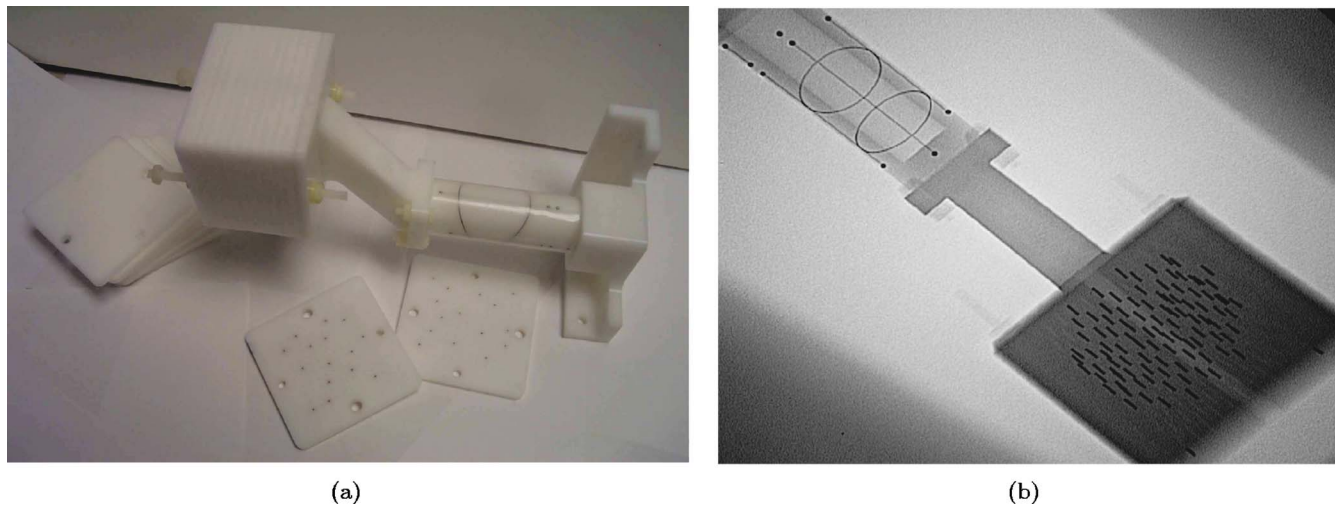
FIG. 9. (a) An image of the seed phantom attached to the FTRAC fiducial. The seed phantom can replicate any implant configuration, using the twelve 5 mm slabs each with over a hundred holes. (b) A typical x-ray image of the combination.

center of the image. Averaged % matching from 27 500 data sets are plotted against the mean image-distortion in Fig. 8. Even with 15 mm distortion errors, % matching is above 95%. Average image distortion of up to 5 mm matches over 99% of the seeds correctly. Thus MARSHAL seems to be fairly robust to moderate distortion errors. Again, it fares well since it solves the matching problem as a global combinatorial optimization. Similar results are also available from phantom experiments in Sec. III D.

C. Phantom experiments

After the simulations, experiments were conducted on a precisely fabricated seed phantom, constructed using acetol. The FTRAC fiducial[25] was used to track the C-arm (accuracy of 0.56 mm translation and 0.33° rotation), and was attached to the seed phantom as shown in Fig. 9. The seed phantom is comprised of twelve 5-mm-thick slabs, each having at least a hundred holes with 5 mm spacing. Any implant configuration with accurately known seed positions (in the FTRAC frame) can be created. Unfortunately there was about 0.5°–1° rotational error in the assembly of the attachment, leading to an error of about 0.5 mm in the ground truth estimates. Thus we expected additional error even before doing the experiments. The phantom is otherwise highly accurate. The seed density was kept constant at about 1.56 seed/cc. The number of seeds (and hence prostate size) was increased from 40 to 100 in steps of 15. For a given constellation, six images within a 20° cone around the AP axis were taken using an *uncalibrated Philips Integris V3000* fluoroscope. The obtained images were dewarped using the pin-cushion test. Accurate ground truth for matching was computed by utilizing the known 3D seed locations.

Matching was achieved using MARSHAL, followed by performance analysis. MARSHAL was evaluated for cases when two, two and a half, three, four, or more images are used. Robustness is further evaluated by using distorted images. For each implant, six images from various poses were

taken, and a number of combinations are created by choosing a different set of images to reconstruct the seeds. The results are averaged when displayed in Tables II and III.

Two images: The results for MARSHAL with two images, evaluated on 75 combinations (15 for each data set), are summarized on the left side of Table II. Two-image matching gives poor results, matching only about 2/3 of the seeds. The matched seeds reconstruct with an average RE of 1.17 mm, while the mismatched seeds reconstruct with an average error of 27.2 mm, the maximum being 241.8 mm. Though these results are completely unacceptable, MARSHAL provides the best possible solution that can exist with two images. Due to the various calibration and numerical errors, alternate constellations have significantly lesser error than the correct one. This is evident from the table where we can see that the average PE for mismatched seeds is only 0.22 mm (0.15 mm RA), but the seeds are deviated by a large value of 27.2 mm in 3D. Thus with two images, MARSHAL produces the best possible (least cost) matching that can exist.

Two and a half images: To *understand* better how information from a third image facilitates seed matching, partial information from a third image was added. Instead of all three, it runs only one bipartite matching, while still using Eq. (3) to compute $\mathcal{C}_{ij}$. Since the third image is not completely used and this exercise is only toward our understanding, we say that only two and a half images are used. The results from 300 combinations ($5 \times {}^6C_2 \times 4$—the choice for the third image has four options) are shown on the right side of Table II. This little information from the third image is sufficient to correctly match 98.3% of the seeds. The mean RE for matched seeds is 0.88 mm, while that of mismatched seeds is 7.68 mm, the maximum for all combinations being 68.2 mm.

Thus when a third image is used, the number of good triplets decreases from $O(N^{5/3})$ to $O(N^{4/3})$, providing a huge improvement in performance. Though the matching rate is

TABLE II. MARSHAL performance on phantom data. Two-image seed matching is completely erroneous (left), with over a quarter of the seeds being mismatched with high reconstruction error, yet low projection error and reconstruction accuracy. Adding just little information from a third image boosts performance (right). Most seeds match, though the reconstruction error is high for the few mismatched seeds. A total of 75 and 300 combinations were used, respectively.

		Number of seeds									
		Two images					Two and a half images				
		40	55	70	85	100	40	55	70	85	100
Matching rate (%)	Match	76.5	75.8	68.3	61.9	53.1	98.3	99.8	98.3	98.0	97.2
Reconstruction error (mm)		0.96	0.92	1.21	1.17	1.59	0.83	0.73	0.88	0.89	1.09
Projection error (mm)		0.19	0.12	0.15	0.15	0.19	0.20	0.16	0.17	0.18	0.24
Reconstruction error (mm)	Mismatch	23.9	27.8	28.2	23.0	33.0	2.35	13.2	4.32	8.38	10.2
Projection error (mm)		0.22	0.19	0.18	0.22	0.27	0.22	0.29	0.16	0.26	0.26
Reconstruction accuracy (mm)		0.15	0.13	0.13	0.15	0.19	0.13	0.09	0.12	0.16	0.15

good, RE for mismatched seeds is still rather high (note that PE for mismatched seeds is low). The reason is that in spite of being very close, the projection of the solution of the tripartite problem is not exactly the same as the solution of the bipartite problem. The proximity in the solution is evident from the low PE of 0.24 mm for the few mismatched seeds ($<3\%$). Thus by including some combinatorial consistency checks, we should be able to match all the seeds. This intuition is realized fully in MARSHAL.

Three images: Averaged results from a total of 100 combinations ($5 \times {}^6C_3$) are shown on the left side of Table III. 98.5% of the seeds matched perfectly. The 3D RE for matched seeds has a mean of 0.63 mm and STD 0.24 mm. The % of mismatched seeds is a strong upper-bound due to

the overlap problem (Sec. II H). This is evident from the RE for mismatched seeds, which has a mean value of 0.91 mm, only a little higher than correctly matched seeds. The average worst case error was 1.32 mm, the error being lower for smaller implants. As mentioned earlier, the ground truth of the seed locations has a small rotation error. To counter this, we computed *relative RE*, which removes any constant translation and rotation offset between the ground truth and reconstructed seeds, and measures any variation in the *shape* of the reconstructed implant. It can be seen that relative RE is significantly lower, standing at a mean value of 0.32 mm. Thus there is a constant shift of about 0.3 mm (due to the error in ground truth) for the whole implant, and an additional error of 0.32 mm for each seed. The absolute worst

TABLE III. Performance on phantom data. Using three images gives excellent results (left), with most of the seeds being matched. Mismatched seeds reconstruct with a low error. Using four images gives similar results (right). A total of 100 and 300 combinations were used, respectively.

		Number of seeds									
		Three images					Four images				
		40	55	70	85	100	40	55	70	85	100
Matching rate (%)	Match	97.6	100	98.0	98.5	98.5	99.2	100	98.7	98.9	99.0
Reconstruction error (mean)		0.60	0.48	0.63	0.70	0.76	0.59	0.48	0.63	0.69	0.75
Reconstruction error (STD)		0.21	0.23	0.25	0.24	0.25	0.21	0.23	0.25	0.24	0.25
Reconstruction error (mean)	Mismatch	0.73	⋯	0.76	0.84	1.30	0.67	⋯	0.65	0.64	1.14
Reconstruction error (worst)		1.18	⋯	1.03	1.10	1.96	0.91	⋯	0.86	0.77	1.58
Reconstruction error (relative)	All	0.28	0.29	0.35	0.30	0.39	0.28	0.29	0.35	0.29	0.38

case error was 3.29 mm, which seems to have a likelihood of about 1 in 75, as indicated from our data.

Four or more images: To seek further improvement, 300 combinations ($5 \times {}^6C_3 \times 3$) using four images were run, the results of which are summarized on the right side of Table III. It can be seen that, on average, 99.2% of the seeds match perfectly. RE for matched seeds still has a mean of 0.63 mm and STD 0.24 mm. While RE of mismatched seeds has decreased to 0.78 mm and the average worst case error to 1.03 mm, the relative RE is still 0.32 mm. Thus we can see that adding a fourth image does not significantly improve the mismatched error, predominantly because three-image matching leaves very little scope for it. It should be noted that due to this redundancy, a fourth image is a great source for validation during the procedure.

Similar to using a fourth image, using a fifth or sixth image has only a little improvement in performance. Average results from 300 combinations indicate that using five images perfectly matches 99.2% of the seeds with the average worst case error dropping to 0.9 mm. Six images images match 99.4% with the average worst case error being 0.85 mm. Thus three images appears to achieve the correct balance between number of images used and accuracy obtained, though it should be noted that using more extra images would decrease RE and might also aid in the hidden seed problem.

It should be noted that %-matching implies a normalization by the number of seeds. Using four images, the 100-seed implant seems to perform similar to the 40-seed implant, matching around 99% of the total seeds. On the contrary, the 100-seed implant mismatches about 1 seed/implant while the 40-seed implant mismatches only 0.3 seeds/implant. In a similar manner, though the 55-seed implant matches 100% of the seeds, it does not imply that all 55-seed implants would always correctly match. Depending on a host of uncontainable factors like errors in tracking, segmentation, noise, etc., an erroneous variation of about 1 seed/implant should typically be expected while studying any pattern.

D. Distorted images

MARSHAL addresses the problem globally as a combinatorial optimization problem, as compared to locally by using only epipolar constraints. Thus it is well-suited to evaluate the correspondences even in the presence of nonlinear effects like image distortion. Well-separated image positions (with distortion) were shown to match 98.5% of the seeds even on a 100-seed implant.[67] To further validate the robustness, we tested it for 700 combinations of distorted C-arm images. These combinations were chosen, so as to not restrict the images in any way, apart from being inside a 25° cone around the AP axis. Thus the images can be very close to each other and not well separated. Using four such images perfectly matching 93.2% of the seeds on an average, performing superior to three images that matched only 86.3%. Five and six images perfectly matched 95.4% and 96.2% of the seeds, respectively. The performance can be further improved by keeping the seed projections closer to the center of the image. It can be observed from Fig. 9 that the seed phantom is not in the center of the image, and thus suffers from a considerable amount of distortion.

Thus, if the prostate size is small (<50 cc), then MARSHAL with three distorted images can match over 97% of the seeds match perfectly. This means that it flips about 2–3 seeds on average, correctly matching the rest. Alternatively, if the prostate is a little larger, then using more images that are well-spaced can still match about 98% of the seeds. Average RE for the whole implant stood high at 2.65 mm, while relative RE was low at 0.54 mm. The relative RE for the mismatched seeds was about 1.4 mm, while the absolute RE was 3.9 mm. Thus most seeds matched correctly, though the reconstructed implant is shifted in space by about 2 mm due to image distortion. The mismatched seeds have an additional error of 1 mm after the shift. Similar patterns were also observed on synthetic data in Sec. III B. Though we have not encountered any similar analysis in the literature, we conjuncture that since those algorithms rely on epipolar lines that are distorted nonlinearly, they should be affected much more severely than MARSHAL.

It should be noted that though theoretically the reconstructed implant should be skewed, practically it is of the same shape. This is evident from the marginal increase in relative RE by only 0.23 mm. Thus the implant reconstructed from distorted images can be practically shifted by about 2 mm to overlap with the correct location. The reason for the absence of skew is that even though distortion is always nonlinear, the area that the implant occupies in the image is small. This results in an approximately linear shift of the seed coordinates in the image. Furthermore, 3D reconstruction from multiple images also reduces this error. As a consequence, the reconstructed implant is practically shifted in space in proportion to the amount of distortion.

Thus image distortion can potentially be either corrected using a constant function, or completely avoided if the amount is low. It should be noted that for this to work the implant should not be large and some mechanism needs to correct for the consistent shift in the x-ray to TRUS registration step. Moreover, additional images can provide significant improvement and reliability.

E. Run-time

MARSHAL has a run-time complexity of $O(N^3)$ for any number of images, which makes it extremely fast even on large size implants. The unoptimized MATLAB code (Pentium4 3.2 GHz, Windows 2000, 1 Gbyte RAM) on a 100-seed implant with a seed density of 1.8 seeds/cc ran in 1.6 s. This run-time is typical of most tested implants, indicating that an optimized C++ implementation could in practice run in near real-time. Thus run-time for MARSHAL seems to be sufficient for intraoperative dose, especially when compared to the time required for automatic segmentation, which typically requires a few seconds. Though it was never observed in any of the thousands of test data sets, we point out that there can exist (theoretically) many long loops, which can potentially slow down MARSHAL.

IV. SHORTCOMINGS AND FUTURE WORK

The underlying assumption of this paper has been that all the seeds are segmented and their image coordinates known. Though excellent segmentation algorithms have been proposed,[15,35,36] they are capable of only segmenting seeds that are fully visible. Thus the coordinates of any hidden seeds (due to an overlap) will not be available, resulting in varying number of segmented seeds in each image. This is an issue of utmost importance and clinical practicality. While some researchers have reported the number of hidden seeds to be between 7% and 45% of the total seeds,[35] others have noticed only about 2%–3%[68] in their clinical data. Nevertheless, no solution can be clinically viable until it explicitly tackles this problem.

In some of the suggested solutions, fast-CARS was extended[35] to tackle this, but the algorithm reconstructed a greater number of seeds than were actually inserted. Another variant of fast-CARS was proposed[68] by ordering the seeds using the epipolar constraints. Unfortunately, the algorithm "*based on epipolar geometry requires co-planar imaging to perform reconstruction*" and "*cannot reconstruct undetected seeds if they exist in the same search restriction band*." Significant works as they are, the problem merits further research. An intensity-based method using tomosynthesis[69] has also been proposed. Unfortunately in their current implementation, they require a rather large number of images to achieve accurate reconstruction. Another method based on Hough trajectories[70] has been proposed to solve this problem, but again requires a large number of images. Nevertheless, these methods could lead to promising results in the future.

Though MARSHAL does not yet tackle hidden and spurious seeds, we believe that the proposed mathematical framework for seed matching is comprehensive. It can be extended to solve these problems in a more formal framework of combinatorial optimization in contrast to the previously mentioned heuristic methodologies. This formal approach is expected to provide a better control on the behavior of the algorithm, potentially offering a decrease in dependence on rules that might change behavior from one data set to another and an increase in robustness. Moreover, the complete problem can be tackled as a whole (global optimization) as compared to looking at one seed or a group of seeds with heuristic rules (local optimization). Though it might be argued that the full problem should be broken down into independent subset of seeds with each subset solved separately, this might not be feasible in the generic case. If the tripartite problem could be broken down into a set of smaller min-cost tripartite problems, then it would no longer be NP-hard. Thus any solution should try to address the problem as a whole.

The basics of our approach will be to redefine the problem into a formal framework. Since we have proved in this publication that the global optimum is not attainable by any polynomial algorithm, we will project the full problem into multiple lower dimensional problems involving all the seeds. Each of these subproblems would be casted into the network flow framework and the global optima for each lower dimensional problems would be achieved in polynomial time. The solutions to the subproblems would later be integrated together to create the final solution. The predominant difference from the current framework will be that instead of forced unitary flows, variable flows will be handled. Thus all the seeds will be handled in unison and solved globally. In addition, theoretically assured bounds on the performance of the algorithm are currently being worked out. A combination of the above-mentioned promises to lead to an implementation that could potentially be used clinically.

V. CONCLUSION

In contrast to previously proposed methods, we have formalized the seed-matching problem. We have proposed an algorithm, abbreviated as MARSHAL. It typically runs with $O(N^3)$ time complexity with any number of images. Using three images, it matched perfectly over 98.5% of the seeds at a density of 1.56 seeds/cc. MARSHAL in combination with the FTRAC fiducial can reconstruct the seeds with a mean error of 0.63 mm and a STD of 0.24 mm. The mismatched seeds had a mean error of 0.9 mm, while the most mismatched seed in each implant had an average error of 1.32 mm. The worst error across all data sets was 3.29 mm and has a low likelihood of 1 in 75 implant cases. Moreover, relative reconstruction error was 0.32 mm, which reflects the deviation in shape of the reconstructed implant when compared to the true implant.

The algorithm performs well for distorted images too, matching over 97% of the seeds for small implants. The simulation experiments indicate that MARSHAL is not sensitive to (a) image separation, (b) seed density, (c) number of seeds, and (d) C-arm calibration. It is also acceptably robust to (a) segmentation, (b) C-arm pose, and (c) distortion. It can reconstruct an implant when three of more images are used, with a robustness, precision, and speed that promises to be sufficient to support intraoperative dosimetry in prostate brachytherapy. Though we validate it only for brachytherapy, MARSHAL is sufficiently generic to be used for establishing correspondences across many synergistic applications.

[a])Electronic mail: jain@cs.jhu.edu
[b])Electronic mail: gabor@cs.jhu.edu
[1]A. Jemal, "Cancer statistics," Ca-Cancer J. Clin. **54**(1), 8–29 (2004).
[2]R. Peschel and J. Colberg, "Surgery, brachytherapy, and external-beam radiotherapy for early prostate cancer," Lancet Oncol. **4**, 233–241 (2003).
[3]J. Blasko, T. Mate, J. Sylvester, P. Grimm, and W. Cavanagh, "Brachytherapy for carcinoma of the prostate: Techniques, patient selection, and clinical outcomes," Semin. Radiat. Oncol. **12**, 81–94 (2002).
[4]G. Merrick, W. Butler, J. Lief, and A. Dorsey, "Is brachytherapy comparable with radical prostatectomy and external-beam radiation for clini-

cally localized prostate cancer?," Tech. Urol. **7**, 12–19 (2001).

[5] S. Nag, J. Ciezki, R. Cormack, S. Doggett, K. DeWyngaert, G. Edmundson, R. Stock, N. Stone, Y. Yu, and M. Zelefsky, "Intraoperative planning and evaluation of permanent prostate brachytherapy: Report of the American Brachytherapy Society," Int. J. Radiat. Oncol., Biol., Phys. **51**, 1422–1430 (2001).

[6] J. Xue, E. Gressen, and T. Jefferson, "Feasibility of trus-based prostate post-implant dosimetry," AAPM Annual Meeting, July 2004, p. Poster.

[7] V. A. Dumane, M. Zaider, G. N. Cohen, and W. M. F. Worman, "Combined ultrasound-fluoroscopy approach to the intraoperative detection of seeds in prostate brachytherapy," ASTRO Annual Meeting, *Poster*, Oct 3–7 2004.

[8] F. Mitri, P. Trompette, and J. Chapelon, "Improving the use of vibro-acoustography for brachytherapy metal seed imaging: A feasibility study," IEEE Trans. Med. Imaging **23**, 1–6 (2004).

[9] D. Holmes, B. Davis, C. Bruce, and R. Robb, "3d visualization, analysis, and treatment of the prostate using trans-urethral ultrasound," Comput. Med. Imaging Graph. **27**, 339–349 (2003).

[10] D. French, J. Morris, M. Keyes, and S. E. Salcudean, "Real-time dosimetry for prostate brachytherapy using trus and fluoroscopy," MICCAI 2004, pp. 983–991.

[11] B. Han, K. Wallner, G. Merrick, W. Butler, S. Sutlief, and J. Sylvester, "Prostate brachytherapy seed identification on post-implant trus images," Med. Phys. **30**, 898–900 (2003).

[12] P. Kumar, R. Good, B. Epstein, M. Hussain, and F. Bartone, "Fluoroscopy guided transperineal percutaneous permanent 125iodine implantation of prostate cancer," Radiat. Med. **3**, 161–167 (1985).

[13] B. Prestidge, J. Prete, T. Buchholz, J. Friedland, R. Stock, P. Grimm, and W. Bice, "A survey of current clinical practice of permanent prostate brachytherapy in the United States," Int. J. Radiat. Oncol., Biol., Phys. **15**, 461–465 (1998).

[14] P. Ravindran, C. Lewis, J. Alphonsi, P. Lindsay, and D. D'Souza, "A c-arm based intra-operative dosimetry system for trans-rectal ultrasound guided prostate implant with 125I seeds," in Proceedings of the 14th International Conference on the Use of Computers in Radiation Therapy, 2004.

[15] D. Todor, M. Zaider, G. Cohen, M. Worman, and M. Zelefsky, "Intraoperative dynamic dosimetry for prostate implants," Phys. Med. Biol. **48**, 1153–1171 (2003).

[16] D. Todor, G. Cohen, H. Amols, and M. Zaider, "Operator-free, film-based 3d seed reconstruction in brachytherapy," Phys. Med. Biol. **47**, 2031–2048 (2002).

[17] L. Gong, P. Cho, B. Han, K. Wallner, S. Sutlief, S. Pathak, D. Haynor, and Y. Kim, "Ultrasonography and fluoroscopic fusion for prostate brachytherapy dosimetry," Int. J. Radiat. Oncol., Biol., Phys. **1**, 1322–1330 (2002).

[18] G. Grado, T. Larson, C. Balch, M. Grado, J. Collins, J. Kriegshauser, G. Swanson, R. Navickis, and M. Wilkes, "Actuarial disease-free survival after prostate cancer brachytherapy using interactive techniques with biplane ultrasound and fluoroscopic guidance," Int. J. Radiat. Oncol., Biol., Phys. **1**, 289–299 (1998).

[19] G. Merrick, W. Butler, A. Dorsey, and H. Walbert, "Prostatic conformal brachytherapy: 125i/103pd postoperative dosimetric analysis," Radiat. Oncol. Invest. **5**, 305–313 (1997).

[20] D. Nori and J. Moni, "Current issues in techniques of prostate brachytherapy," Semin Surg. Oncol. **13**, 444–453 (1997).

[21] S. Nag, D. Scaperoth, R. Badalament, S. Hall, and J. Burgers, "Transperineal palladium 103 prostate brachytherapy: Analysis of morbidity and seed migration," Urology **45**, 87–92 (1995).

[22] K. Wallner, S. Chiu-Tsao, J. Roy, V. Arterbery, W. Whitmore, S. Jain, B. Minsky, P. Russo, and Z. Fuks, "An improved method for computerized tomography-planned transperineal 125iodine prostate implants," J. Urol. (Baltimore) **146**, 90–95 (1991).

[23] R. Hofstetter, M. Slomczykowski, M. Sati, and L. Nolte, "Fluoroscopy as an imaging means for computer-assisted surgical navigation," Comput. Aided Surg. **4**, 65–76 (1999).

[24] J. Yao, R. H. Taylor, R. P. Goldberg, R. Kumar, A. Bzostek, V. R. Van, P. Kazandizes, and A. Gueziec, "A c-arm fluoroscopy-guided progressive cut refinement strategy using a surgical robot," Comput. Aided Surg. **5**, 373–390 (2000).

[25] A. Jain, T. Mustufa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (ftrac) fiducial," Med. Phys. **32**, 3185–3198 (2005).

[26] A. Fung, "C-arm imaging for brachytherapy source reconstruction: Geometrical accuracy," Med. Phys. **29**, 724–726 (2002).

[27] S. Narayanan and P. Cho, "3d seed reconstruction from unknown imaging geometry," in Proceedings of the 14th International Conference on the Use of Computers in Radiation Therapy, 2004, pp. 762–765.

[28] M. Murphy and D. Todor, "Demonstration of a forward iterative method to reconstruct brachytherapy seed configurations from x-ray projections," Phys. Med. Biol. **50**, 2715–2737 (2005).

[29] A. Jain, R. Kyon, and G. Fichtinger, "C-arm calibration—Is it really necessary?," Medical Image Computing and Computer Assisted Interventions (MICCAI), 2005.

[30] K. Wallner, J. Roy, M. Zelefsky, Z. Fuks, and L. Harrison, "Fluoroscopic visualization of the prostatic urethra to guide transperineal prostate implantation," Int. J. Radiat. Oncol., Biol., Phys. **1**, 863–867 (1994).

[31] P. Roberson, V. Narayana, D. McShan, R. Winfield, and P. McLaughlin, "Source placement error for permanent implant of the prostate," Med. Phys. **24**, 251–257 (1997).

[32] M. Baird, R. Holt, and T. Selby, "Improvement of transperineal implant dosimetry by intraoperative cystoscopic confirmation of prostate anatomy," J. Urol. (Baltimore) **164**, 406–410 (2000).

[33] L. Archambault, L. Beaulieu, and D. Tubic, "Automatic post-implant needle reconstruction algorithm to characterize and improve implant robustness analyses," Med. Phys. **30**, 2897–2903 (2003).

[34] M. Zhang, M. Zaider, M. Worman, and G. Cohen, "On the question of 3d seed reconstruction in prostate brachytherapy: The determination of x-ray source and film locations," Phys. Med. Biol. **49**, N335–345 (2004).

[35] Y. Su, B. Davis, M. Herman, and R. Ra, "Prostate brachytherapy seed localization by analysis of multiple projections: Identifying and addressing the seed overlap problem," Med. Phys. **31**, 1277–1287 (2004).

[36] D. Tubic, A. Zaccarin, J. Pouliot, and L. Beaulieu, "Automated seed detection and three-dimensional reconstruction. II. Reconstruction of permanent prostate implants using simulated annealing," Med. Phys. **28**, 2272–2279 (2001).

[37] S. Lam, P. Cho, and R. Marks, "Prostate brachytherapy seed segmentation using spoke transform," Proc. SPIE **4322**, 1490–1500 (2001).

[38] P. Cho, "Computerized segmentation of clustered seeds in prostate brachytherapy," International Conference on the Use of Computers in Radiation Therapy, 2000, pp. 105–107.

[39] H. Amols and I. Rosen, "A three-film technique for reconstruction of radioactive seed implants," Med. Phys. **8**, 210–214 (1981).

[40] P. Biggs and D. Kelley, "Geometric reconstruction of seed implants using a three-film technique," Med. Phys. **10**, 701–704 (1983).

[41] M. Rosenthal and R. Nath, "An automatic seed identification technique for interstitial implants using three isocentric radiographs," Med. Phys. **10**, 475–479 (1983).

[42] M. Altschuler, P. Findlay, and R. Epperson, "Rapid accurate, three-dimensional location of multiple sees in implant radiotherapy treatment planning," Phys. Med. Biol. **28**, 1305–1308 (1983).

[43] D. Jackson, "An automatic method for localizing radioactive seeds in implant dosimetry," Med. Phys. **10**, 370–372 (1983).

[44] S. Li, G. Chen, C. Pelizzari, C. Reft, J. Roeske, and Y. Lu, "A new source localization algorithm with no requirement of one-to-one source correspondence between biplane radiographs," Med. Phys. **23**, 921–927 (1996).

[45] R. L. Siddon and L. M. Chin, "Two-film brachytherapy reconstruction algorithm," Med. Phys. **12**, 77–83 (1985).

[46] M. Altschuler and A. Kassaee, "Automated matching of corresponding seed images of three simulator radiographs to allow 3d triangulation of implanted seeds," Phys. Med. Biol. **42**, 293–302 (1997).

[47] S. Narayanan, P. Cho, and R. Marks, "Fast cross-projection algorithm for reconstruction of seeds in prostate brachytherapy," Med. Phys. **29**, 1572–1579 (2002).

[48] K. Thornton, "Prostate brachytherapy seed localization by fluoroscopy/ultrasound fusion: Algorithms and analysis," AAPM Annual Meeting, July 2004.

[49] H. Liu, G. Cheng, Y. Yu, R. Brasacchio, D. Rubens, J. Strang, L. Liao, and E. Messing, "Automatic localization of implanted seeds from post-implant ct images," Phys. Med. Biol. **48**, 1191–1203 (2003).

[50] D. Dubois, W. J. Bice, and B. Prestige, "Ct and mri derived source localization error in a custom prostate phantom using automated image coregistration," Med. Phys. **28**, 2280–2284 (2001).

[51] K. Shafique and M. Shah, "A non-iterative greedy algorithm for multi-frame point correspondence," ICCV, 2003, p. 6.

[52]C. J. Veenman, M. J. T. Reinders, and E. Backer, "Resolving motion correspondence for densely moving points," IEEE Trans. Pattern Anal. Mach. Intell. **23**, 54–72 (2001). [Online.] Available: citeseer.ist.psu.edu/veenman01resolving.html

[53]C. Veenman, M. Reinders, and E. Backer, "Establishing motion correspondence using extended temporal scope," Artificial Intelligience **145**, 227–243 (2003).

[54]V. Salari and I. Sethi, "Feature point correspondence in the presence of occlusion," Pattern Anal. Machine Intell. **12**, 87–91 (1990).

[55]Y. Cheng, R. Collins, A. Hanson, and E. Riseman, "Triangulation without correspondences," ARPA94, 1994, pp. II:993–1000. [Online.] Available: citeseer.ist.psu.edu/6193.html

[56]Y. Cheng, V. Wu, R. Collins, A. Hanson, and E. Riseman, "Maximum-weight bipartite matching technique and its application in image feature matching," 1996. [Online.] Available: citeseer.ist.psu.edu/cheng96maximumweight.html

[57]*Artificial Intelligence: A Modern Approach*, Prentice Hall Series in Artificial Intelligence, edited by S. Russell and P. Norvig, 2nd ed., Chap. 24.5, pp. 885–892.

[58]T. H. Cormen, C. E. Leiserson, and R. Rivest, *Introduction to Algorithms* 2nd ed. (MIT Press, 1990).

[59]J. Edmonds and R. M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," J. ACM **19**, 248–264 (1972).

[60]A. Goldberg, "Network optimization library," http://www.avglab.com/andrew/soft.html.

[61]Algorithmic Solutions Software GmbH, "Library of efficient data types and algorithms (leda)," www.algorithmic-solutions.info/leda_guide/graph_algorithms/mincost_flow.html.

[62]R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, 1st ed. (Prentice Hall, Englewood Cliffs, NJ, 1993).

[63]H. W. Kuhn, "The Hungarian method for the assignment problem," Naval Res. Logistics Quart. **2**, 83–97 (1955).

[64]C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Prentice-Hall, Englewood Cliffs, NJ, 1982).

[65]G. Herman and A. Kuba, *Discrete Tomography: Foundations, Algorithms, and Applications* (Birkhauser, Boston, 1999).

[66]K. Wallner, J. Blasko, and M. Dattoli, *Prostate Brachytherapy Made Complicated*, 2nd ed. (Smart Medicine, 2001).

[67]A. Jain, Y. Zhou, T. Mustufa, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)," SPIE Medical Imaging; Visualization, Image-Guided Procedures, and Display, 2005.

[68]S. Narayanan, P. Cho, and R. Marks, "Three-dimensional seed reconstruction from an incomplete data set for prostate brachytherapy," Phys. Med. Biol. **49**, 3483–3494 (2004).

[69]I. Tutar, R. Managuli, V. Shamdasani, P. Cho, S. Pathak, and Y. Kim, "Tomosynthesis-based localization of radioactive seeds in prostate brachytherapy," Med. Phys. **30**, 3135–3142 (2003).

[70]S. Lam, P. Cho, R. Marks, and S. Narayanan, "Three-dimensional seed reconstruction for prostate brachytherapy using hough trajectories," Phys. Med. Biol. **49**, 557–569 (2004).

FTRAC—A robust fluoroscope tracking fiducial

Ameet Kumar Jain[a)]
Department of Computer, Science, Johns Hopkins University, Baltimore, Maryland 21218

Tabish Mustafa
Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218

Yu Zhou
Department of Mechanical Engineering, Stony Brook University, Stony Brook, New York 11794

Clif Burdette
Accoustic MedSystems, Inc., 2110 Clearlake Blvd., Champaign, Illinois 61822

Gregory S. Chirikjian
Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218

Gabor Fichtinger[a)]
Department of {Computer Science, Mechanical Engineering, Radiology}, Johns Hopkins University, Baltimore, Maryland 21218

C-arm fluoroscopy is ubiquitous in contemporary surgery, but it lacks the ability to accurately reconstruct three-dimensional (3D) information. A major obstacle in fluoroscopic reconstruction is discerning the pose of the x-ray image, in 3D space. Optical/magnetic trackers tend to be prohibitively expensive, intrusive and cumbersome in many applications. We present single-image-based fluoroscope tracking (FTRAC) with the use of an external radiographic fiducial consisting of a mathematically optimized set of ellipses, lines, and points. This is an improvement over contemporary fiducials, which use only points. The fiducial encodes six degrees of freedom in a single image by creating a unique view from any direction. A nonlinear optimizer can rapidly compute the pose of the fiducial using this image. The current embodiment has salient attributes: small dimensions ($3 \times 3 \times 5$ cm); need not be close to the anatomy of interest; and accurately segmentable. We tested the fiducial and the pose recovery method on synthetic data and also experimentally on a precisely machined mechanical phantom. Pose recovery in phantom experiments had an accuracy of 0.56 mm in translation and 0.33° in orientation. Object reconstruction had a mean error of 0.53 mm with 0.16 mm STD. The method offers accuracies similar to commercial tracking systems, and appears to be sufficiently robust for intraoperative quantitative *C*-arm fluoroscopy. Simulation experiments indicate that the size can be further reduced to $1 \times 1 \times 2$ cm, with only a marginal drop in accuracy. © *2005 American Association of Physicists in Medicine*. [DOI: 10.1118/1.2047782]

Key words: *C*-arm, fluoroscopy, tracking, reconstruction, registration, prostate brachytherapy

I. INTRODUCTION

C-arm fluoroscopy is the most widely used intraoperative imaging modality in general surgery (some of the prominent works being Refs. 1–8), but it presently lacks the ability for robust and easy quantitative guidance.[9] Quantitative fluoroscopy-guided surgery needs to solve four major problems: (1) *C*-arm image distortion; (2) the calibration of imaging parameters; (3) pose recovery or tracking; and (4) registration to imaging modalities. The first two are well-studied problems in the literature.[1,3,9] On the other hand, pose recovery on unencoded *C*-arm machines is a major technical problem that presently does not have a clinically practical solution in many areas of application. In this paper, we propose a solution to accurately estimate the pose by using a radio-opaque fiducial. Moreover, it can also register the x-ray images to other imaging modalities.

The history of using radio-opaque fiducials to solve for the *C*-arm pose starts as early as 1987,[10] where 16 spherical beads in a known configuration were introduced in the region of interest. Theorems from projection geometry were used to reconstruct the *C*-arm pose, calibration parameters and target details. This was later followed by other bead-based fiducials[2,11–13] for solving both the problems of a *C*-arm pose and calibration (a small survey[14]). The number of beads varied from 20 to 613, with the patient sometimes fitting inside them. Though the accuracies were acceptable (0.5 – 1 mm), the fiducials tended to be (a) too large; (b) cumbersome to use in a clinical setting; (c) interfering with the anatomy in the image; and (d) nontrivial to segment. A smaller and well-encoded fiducial can solve these problems efficiently.

Off-the-shelf tracking devices, as they became available, were a natural choice to simplify the problem. In current commercial *C*-arm fluoroscopy surgical navigation systems,

the x-ray detector is localized in room coordinates by some auxiliary optical tracker[15,16] or electromagnetic (EM) tracker.[17] Unfortunately, auxiliary trackers sometimes become impractical for various reasons. They are expensive and add to the complexity of the operating room since they require an additional calibration step. Optical trackers require a line of sight, which becomes cumbersome in a clinical setting and requires an alteration in the standard workflow. The EM trackers can successfully overcome this issue, but become susceptible to field distortion from metal objects like surgical tools or the *C* arm itself, compromising on accuracy. This has lead some recent researchers to prefer fiducial-based tracking. In a recent publication,[18] the authors delineate the above problems and also say that using optical trackers reduces the useful imaging volume of the fluoroscope and potentially compromises the achievable accuracies. Despite using an optical tracker to track their surgical tools, they explicitly choose to not track the *C* arm using the tracker but instead use a radio-opaque fiducial. Their system has been fairly successful for various surgeries and has been in clinical use for the last four years. Similar preferences can be observed in other recent publications as well.[3,19–22]

To make the fiducials feasible, recent publications have reported smaller fiducials by compact bead placement. It should be noted that this increase in clinical friendliness by decreasing the size was achieved at the cost of a decrease in accuracy. The typical number of beads were reported to be between 6–28,[18–22] achieving translation accuracies ranging from $1-3$ mm and orientation accuracies around $1°–2°$. The variation in accuracies is governed by both bead configuration[19] and implementation[20] choices. Increasing the number of beads greater than 6 provides little improvement in the accuracy.[19] Moreover, nonlinear optimization fares far superior when compared to linear methods.[20,23] Thus, from recent publications, it seems that a 1 mm error in translation and 1° error in rotation is probably the best that bead-based fiducials can achieve. The reason for this is that accurate segmentation of beads in the x-ray image is nontrivial. A 2 mm diameter radio-opaque bead would project as an ellipse with a 12 pixel long axis. The center of the ellipse does not necessarily correspond to the center of bead, contributing at least $1-2$ pixel segmentation error. Moreover, a further decrease in fiducial size would greatly compromise the accuracy.[24]

Methods to accommodate for and decrease the effect of such systematic biases[25] are studied by the computer vision community during their camera calibration procedures. Unfortunately, the improvement is not substantial. Moreover, point-based registration methods are known to be sensitive to segmentation errors.[24] Thus a totally new approach toward fiducial design will be required. Moreover, automatic segmentation and establishing correspondences between the fiducial beads and their projections in the image is itself a nontrivial task for arbitrary positions of the fiducial. In a synergistic problem, the end effector of a surgical robot was registered to a *C* arm using line fiducials[3] mounted on the robot's tool holder. The idea was to register the x-ray image to a robot, in order to locate the target anatomy in the coordinate frame of the surgical robot. Although their helical fiducial did not allow for sufficiently accurate and robust pose recovery, it was an important effort that inspired our work. Our present contribution to the state of the art is the development of a novel fiducial system and its mathematical framework. The fiducial uses ellipses and straight lines in addition to points. It offers a decrease in size, an enhancement in accuracy and robustness in *C*-arm pose recovery, allowing for precise object reconstruction. Some preliminary studies have been published earlier,[21,26] while we present the complete mathematical framework, results and analysis in this paper.

II. MATERIALS AND METHODS

Our solution to image-based pose recovery is to mount a stationary fiducial in the field of the *C*-arm fluoroscope. The fiducial, by design, encodes six degrees of freedom (DOF) from a single image by creating a unique view from any direction. The accuracy and robustness of pose recovery would critically depend on the design and precise manufacturing of the fiducial. Software was created for segmentation of the fiducial and for numerical optimization that determines the six-DOF transformation between the coordinate frames of the *C* arm and the fiducial.

A. Fiducial design and manufacturing

For practical viability the fiducial should (a) be nonintrusive with small dimensions; (b) not corrupt image quality; (c) not necessarily be close to the anatomy of interest; (d) be easily/accurately segmentable from x-ray images; and (e) have software that is robust, fast, and simple to use.

1. Choice of fiducial features

Geometric aspects of x-ray imaging are modeled using the perspective projection model (Sec. II B), under which we can limit the choice of features to well-behaved features like points, straight lines, ellipses, parabolas, helixes, and other curves. Different features exhibit varying segmentation accuracies, each combination thus producing a fiducial with different characteristics. Segmentation is crucial, since it directly relates to achievable pose estimation accuracy. Image processing algorithms can more accurately segment constrained (parametric) curves when compared to unconstrained curves, usually to a subpixel precision.

Projected point features have a high segmentation error[25] ($\sim1-2$ pixel) but make the pose estimation problem very well constrained.[27] Straight lines are not as well constrained as points, but offer superior segmentation accuracies to points. Linear methods are available for both features,[28–30] but it should be noted that due to an inability to incorporate all the rotation constraints, linear methods have been shown to not fare as well as nonlinear methods.[20,23] A helix is a parametric curve that projects as a nonparametric curve with nondifferentiable segments, leading to inaccurate segmentation, which makes pose estimation prone to errors.[3] Alternatively, ellipses, parabolas, and hyperbolas are conics that will
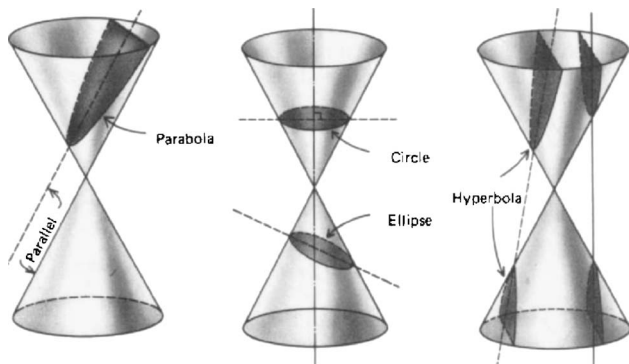
FIG. 1. The center of the cone is the x-ray source, while the base is the image plane. The intersection of this plane with the cone is an ellipse. Any 3D curve combined with the x-ray source defines a unique cone, which intersects the image plane as an ellipse. Any conic in general projects as another conic under perspective geometry.

always project as another member of the same family. The proof of this is self-explanatory from Fig. 1. Moreover, since the mean direction of the x-ray beam is near-orthogonal to the image plane, we can, without loss of generality, assume that the projection of all conics would be an ellipse on the image plane. Parabolas and hyperbolas would project as open ellipses, making segmentation prone to error. In comparison, a 3D ellipse will project as a closed ellipse in the image, making it an ideal candidate feature.[31,32] A 3D circle is a special case of the ellipse.[33–35] It should be noted here that even though ellipses are easy to segment, exact linear techniques for pose estimation from ellipses are not available.[36,37] Some authors have attempted to linearize pose estimation using an ellipse,[29] but not with high accuracy.

Thus we choose a combination of ellipses, straight lines, and points as our primary features, balancing between accuracy and convergence. Ellipses make the algorithm more accurate, points make it more robust, and lines will achieve a bit of both.

2. Necessary and sufficient feature combinations

A combination of the above features can be used for pose estimation. The number and placement of the features will decide the robustness of the fiducial. It can be shown geometrically that the correspondence of at least three points is required for determining the pose of the fiducial uniquely. Using the point-line duality, we can also conclude that a minimum of three straight lines (not line segments but infinite lines) is necessary to find the pose uniquely. A single 3D circle limits the pose to four discrete symmetric choices, along with an indeterminable rotation along the normal axis. An ellipse is able to limit this to only 16 discrete positions. Thus, an ellipse/circle along with a point, can uniquely encode the pose of the fiducial. Open curves like the helix, parabola, and hyperbola can limit the pose to just two discrete rotations.

Note that these are purely theoretical considerations and using the bare minimum number of features will result in a loss of accuracy. Inaccuracies in the imaging model coupled
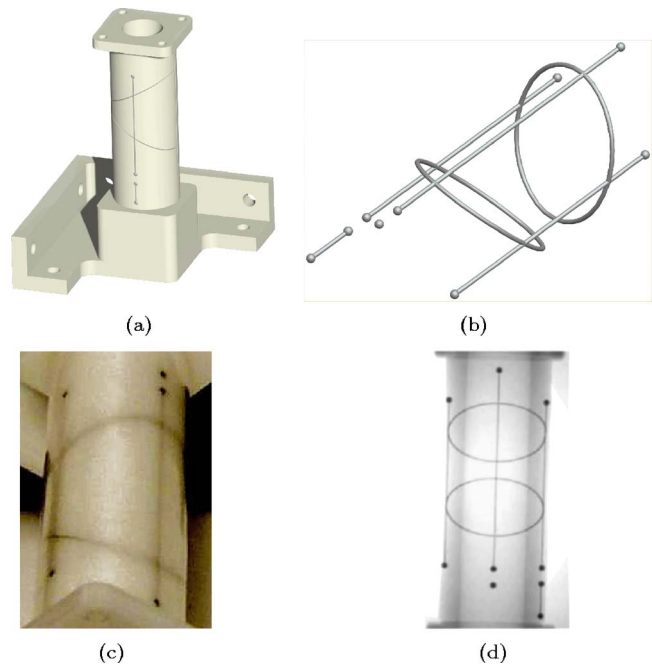

FIG. 2. Images of the FTRAC fiducial (a) cylindrical sheath design; (b) wire model; (c) photograph; (d) x-ray image.

with image quantization errors, would make slight variations from the actual pose produce nearly identical images. In other terms, the Jacobian for image formation loses rank and becomes singular. Singular matrices are undesirable because they cannot be inverted, and indicate that the system has become degenerate. Each feature has blind spots, near which the Jacobian would be singular. Near-singular matrices are theoretically invertible but numerically unstable, resulting in large errors during matrix inversion. For example, the ellipse projects as a straight line in its own plane, and hence the Jacobian would be ill conditioned (near-singular) when the actual pose is in the vicinity of this blind spot. Thus, the placement of all the features should be such that a singularity from one feature is cancelled by a good Jacobian from another. This suggests that a highly redundant structure is needed to accommodate a sufficient range of projection angles and compensate for inevitable calibration and image processing errors.

Our first motivating application is prostate brachytherapy (Sec. IV), where the fluoroscope has a limited view within a 25° cone around the AP axis. To tackle blind spot singularity, we choose to arrange two noncoplanar ellipses 60° from each other (Fig. 2). When one ellipse projects with high eccentricity, the other ellipse projects as a circle, forcing at least one ellipse to encode strongly in any direction. We added three line segments as features, further utilizing the endpoints of the segments as six point features. To make it nonsymmetric, three additional points were added at known distances along the lines. Thus, one segment has four points, another has three, while the third has just two. This fact is also used to distinguish between the ellipses, which otherwise have the same dimensions. These features constitute our fiducial, which henceforth we shall refer to as the fluoroscope track-
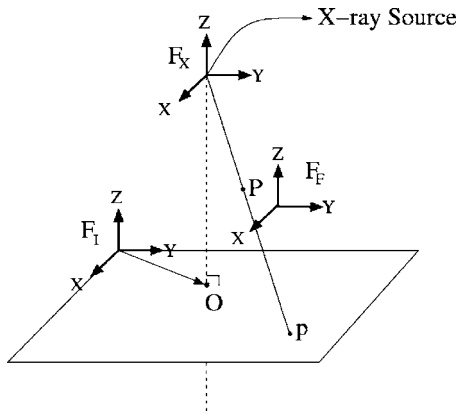
FIG. 3. Projective geometry and notations for fluoroscopic imaging.

ing (FTRAC) fiducial. The features are placed carefully so that the projection of radioactive seeds in the prostate does not overlap with that of the fiducial, even when they are beside each other. All the features of this fiducial are mounted on the surface of a hollow cylinder, which remains stationary. The diameter of the cylinder is 3 cm. Thus, all the features fit inside a $3 \times 3 \times 5$ cm volume.

3. Manufacturing

The design for the fiducial was refined by generating multiple rough prototypes from ABS (acrylonitrile butadiene styrene) using an FDM (fused deposition modeling) rapid prototype machine. The final fiducial design was then fabricated from a acetol rod using a four-axis CNC (computer numerical control) mill, with the grooves of the fiducial geometry machined into its surface to press fit 0.5 mm stainless steel wires and 2 mm stainless steel beads. The cylinder was press fit into a custom acetol mount that provided three mutually orthogonal sets of mounting holes for attaching to an accurate rotary table for validation.

B. Pose recovery mathematics

1. C-arm imaging

Geometric aspects of fluoroscopic imaging can be modeled as a perspective transformation with five parameters—focal length, image origin, and pixel size (Fig. 3). The transformation formula is given in Eqs. (1) and (2); where F_F, F_X, F_I are the coordinate frames of the fiducial, x-ray source, and the image, respectively; P is a 3D point; P_F are the homogenous coordinates of P in F_F; p is the projection of P on the image plane; p_I are the homogenous coordinates (in pixels) of p in frame F_I; $^X F_F$ is the 4×4 rigid transformation matrix that transforms a point in F_F to F_X; $^I F_X$ is the 3×4 perspective projection matrix; f is the focal length; $O = (o_x, o_y)$ is the projection of the x-ray source on the image plane (later referred to as the origin); s_x and s_y are the pixel sizes along the X and Y axes of the image,

$$p_I = {}^I F_X \, {}^X F_F \, P_F,\qquad(1)$$

$$
{}^I F_X = \begin{bmatrix} -f/s_x & 0 & o_x & 0 \\ 0 & -f/s_y & o_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix},
$$

$$(2)$$

$$
{}^X F_F = \begin{bmatrix} r_{11} & r_{12} & r_{13} & T_1 \\ r_{21} & r_{22} & r_{23} & T_2 \\ r_{31} & r_{32} & r_{33} & T_3 \\ 0 & 0 & 0 & 1 \end{bmatrix}.
$$

Assuming that the C arm is calibrated ($^I F_X$ is known), the problem remains to estimate the pose of the fiducial ($^X F_F$), given the image. $^X F_F$ consists of six independent parameters—T_1, T_2, T_3 are the translation parameters, while ϕ_1, ϕ_2, ϕ_3 are the three Euler rotation angles along the x, y, and z axes [which determine r_{ij} in Eq. (2) uniquely].

In this framework, the exact solution is rendered nonlinear and we use the iterative Newton's optimization method to solve it. We segment the fiducial from the image, solve the correspondence problem using the nonsymmetry of the design, and then feed this to the optimizer. The optimizer as a first step computes an expected image location of the fiducial. It also computes a Euclidean-distance-based multidimensional error function. It is assumed that at the correct pose, this error is zero. A closed form of the Jacobian is provided, to ensure quick and accurate convergence.

2. Error formulas for points, lines, and ellipse

The multidimensional error vector measures the Euclidian distance (in the image) between the segmentation of the observed projection of the fiducial, and that of the current projection. A segmentation algorithm for the current image is not needed since the position of any feature in the current image is computed directly using its 3D location on the fiducial. The distance from each feature in the current projection, to its corresponding counterpart in the observed image, constitutes one element in the vector. When all features match perfectly, each element in the vector is zero.

The Euclidian distance for corresponding point features is the difference in their image coordinates. Let (X_i^m, Y_i^m, Z_i^m) be the coordinate of P_i in F_F (model/FTRAC frame) and (X_i, Y_i, Z_i) be its coordinates in F_X. Let (x_i, y_i) be the current coordinates of p_i in F_I. p_i is a function of the six parameters $(\phi_1, \phi_2, \phi_3, T_1, T_2, T_3)$, which we shall refer to by a rotation matrix $\mathcal{R}$ and a translation vector $\mathcal{T}$. Let $(\bar{x}_i, \bar{y}_i)$ be the coordinates of the observed point p_i in the image. Since each iteration of the method has its own rotation and translation parameters, we shall refer to them by their iteration number k as, $\mathcal{R}_k$ and $\mathcal{T}_k$, with $\mathcal{R}_0$ and $\mathcal{T}_0$ being the initial estimate. δ is the error function. Under this notation, the current estimate of the distance function for points can be written as

$$\delta x_i = x_i(\mathcal{R}_k, \mathcal{T}_k) - \bar{x}_i,\qquad(3)$$

$$\delta y_i = y_i(\mathcal{R}_k, \mathcal{T}_k) - \bar{y}_i.\qquad(4)$$

Note that the L_2 norm of these two values is the exact Euclidian distance between the two points, and that the optimizer minimizes the L_2 norm of the vector.

There is no notion of the Euclidian distance between two straight lines. Thus, we select a point on one straight line and measure the distance of this point from the other straight line. Any two points uniquely determine a line, and hence they can be sufficient to measure the distance between the straight lines. If the equation of a straight line in the image is $Ax+By+C=0$, then the distance of (x_i,y_i) from the line is

$$\delta(x_i,y_i) = \frac{Ax_i + By_i + C}{\sqrt{A^2 + B^2}}. \tag{5}$$

Any two points on the line segment can be used, but we use the endpoints for robustness. A similar technique can be used for an ellipse, but no simple closed-form formula exists for the Euclidian distance of a generic point from an ellipse. In fact, this distance is one of the solutions to a fourth degree polynomial,[38] from which the minimum distance solution has to be chosen. We derive an approximation that works well when the point is not too far away from the ellipse. The general equation of an ellipse (conic) is $P^TMP=0$,

$$[x \quad y \quad 1]\begin{bmatrix} A & B/2 & D/2 \\ B/2 & C & E/2 \\ D/2 & E/2 & F \end{bmatrix}\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0, \tag{6}$$

which can be rotated, translated, and scaled to give $P^TF^TMFP=0(P^TM'P=0)$,

$$[x \quad y \quad 1]\begin{bmatrix} 1/a^2 & 0 & 0 \\ 0 & 1/b^2 & 0 \\ 0 & 0 & -1 \end{bmatrix}\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0, \tag{7}$$

where a,b are the length of the major and minor axes. P^TMP is a good distance metric since it is rotation and translation invariant. Unfortunately, it is a nonlinear metric but can be scaled near the ellipse boundary to give an approximate Euclidian distance. Using Eq. (7), the scale factor after some first-order approximations turns out to be $K=ab(AC-B^2/4)/2\sqrt{a^2+b^2}\det(M)$. The final distance function from any point to the ellipse is shown in Eq. (8). Theoretically, a minimum of 5 points are required to localize an ellipse exactly, though we use 12 for robustness,

$$\delta(x_i,y_i) = KP^TMP$$

$$= \frac{ab}{2\sqrt{a^2+b^2}} \frac{(AC-B^2/4)}{det(M)}[x_i \quad y_i \quad 1]$$

$$\times \begin{bmatrix} A & B/2 & D/2 \\ B/2 & C & E/2 \\ D/2 & E/2 & F \end{bmatrix}\begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix}. \tag{8}$$

3. Jacobian calculation

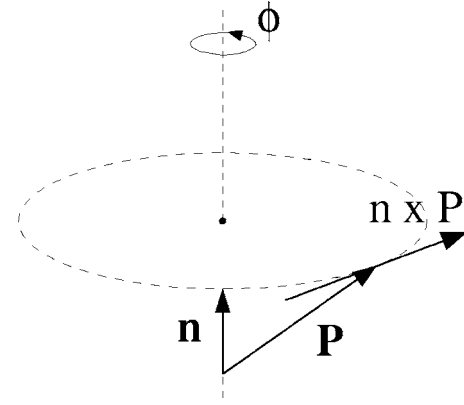The error/distance metric is a 48-dimensional vector (18 from 9 points, 6 from 3 lines, and 24 from 2 ellipses), the

FIG. 4. The derivative of a point vector **P** with respect to a direction **n** is given by the cross-product of **n** and **P**.

zero value of which is computed using the Newton's method. To quickly find this global minimum is a practical challenge. Exact knowledge of the closed form Jacobian of this vector makes a significant improvement in speed and accuracy. The Jacobian is a matrix of size 48×6, and of the form $\partial\delta_i/\partial\mathcal{P}_j$, where $i=1\cdots48$, $j=1\cdots6$; and $\mathcal{P}_j$ representing the (six) pose parameters. We first compute the building blocks $\partial x_i/\partial\mathcal{P}_j$ and $\partial y_i/\partial\mathcal{P}_j$, for later use in the Jacobian computation. From Eqs. (1) and (2), we derive

$$\begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} \dfrac{-fX_i}{s_xZ_i} + o_x \\ \dfrac{-fY_i}{s_xZ_i} + o_y \end{bmatrix}$$

$$= \begin{bmatrix} -\dfrac{f}{s_x}\dfrac{r_{11}X_i^m + r_{12}Y_i^m + r_{13}Z_i^m + T_1}{r_{31}X_i^m + r_{32}Y_i^m + r_{33}Z_i^m + T_3} + o_x \\ -\dfrac{f}{s_y}\dfrac{r_{21}X_i^m + r_{22}Y_i^m + r_{23}Z_i^m + T_1}{r_{31}X_i^m + r_{32}Y_i^m + r_{33}Z_i^m + T_3} + o_y \end{bmatrix}. \tag{9}$$

From here, straightforward manipulations lead to the transnational Jacobian,

$$\frac{\partial x_i}{\partial T_1} = \frac{-f}{s_x}\frac{1}{Z_i}, \quad \frac{\partial x_i}{\partial T_2} = 0, \quad \frac{\partial x_i}{\partial T_3} = \frac{fX_i}{s_xZ_i^2},$$

$$\frac{\partial y_i}{\partial T_1} = 0, \quad \frac{\partial y_i}{\partial T_2} = \frac{-f}{s_y}\frac{1}{Z_i}, \quad \frac{\partial y_i}{\partial T_3} = \frac{fY_i}{s_yZ_i^2}. \tag{10}$$

To obtain the rotational Jacobian, we first estimate $\partial\mathbf{P}_i^m/\partial\phi_j$, which we use to estimate $\partial\mathbf{P}_i/\partial\phi_j$, which will give us $\partial x_i/\partial\phi_j$ and $\partial y_i/\partial\phi_j$. Here $\mathbf{P}_i$ is treated as a vector in space $\partial\mathbf{P}_i^m/\partial\phi_j$ can be obtained geometrically as a cross-product, as illustrated in Fig. 4,

$$\frac{\partial \mathbf{P}_i^m}{\partial \phi_1} = [1,0,0]^T \times \mathbf{P}_i^m,$$

$$\frac{\partial \mathbf{P}_i^m}{\partial \phi_2} = [0,1,0]^T \times \mathbf{P}_i^m, \tag{11}$$

$$\frac{\partial \mathbf{P}_i^m}{\partial \phi_3} = [0,0,1]^T \times \mathbf{P}_i^m.$$

$\partial \mathbf{P}_i^m / \partial \phi_j$ is defined in F_F. It needs to be rotated to change it to frame F_X,

$$\frac{\partial \mathbf{P}_i}{\partial \phi_j} = \mathcal{R} \frac{\partial \mathbf{P}_i^m}{\partial \phi_j}. \tag{12}$$

From Eq. (9), $\partial x_i / \partial \phi_j$ and $\partial y_i / \partial \phi_j$ can be derived as

$$\frac{\partial x_i}{\partial \phi_j} = \frac{-f\left(Z_i \dfrac{\partial X_i}{\partial \phi_j} - X_i \dfrac{\partial Z_i}{\partial \phi_j}\right)}{s_x Z_i^2}, \quad \frac{\partial y_i}{\partial \phi_j} = \frac{-f\left(Z_i \dfrac{\partial Y_i}{\partial \phi_j} - Y_i \dfrac{\partial Z_i}{\partial \phi_j}\right)}{s_y Z_i^2}. \tag{13}$$

Using Eqs. (3)–(5), (8), (10), (12), and (13), we can derive the Jacobian for the point, line, and ellipse features as

$$\frac{\partial (\delta x_i)}{\partial \mathcal{P}_j} = \frac{\partial x_i}{\partial \mathcal{P}_j}, \quad \frac{\partial (\delta y_i)}{\partial \mathcal{P}_j} = \frac{\partial y_i}{\partial \mathcal{P}_j},$$

$$\frac{\partial \delta_j}{\partial \mathcal{P}_j} = \frac{1}{\sqrt{A^2 + B^2}}\left(A \frac{\partial x_i}{\partial \mathcal{P}_j} + B \frac{\partial y_i}{\partial \mathcal{P}_j}\right), \tag{14}$$

$$\frac{\partial \delta_i}{\partial \mathcal{P}_j} = K\left((2Ax_i + By_i + D)\frac{\partial x_i}{\partial \mathcal{P}_j} + (2Cy_i + Bx_i + E)\frac{\partial y_i}{\partial \mathcal{P}_j}\right).$$

4. System of equations and update formula

Thus, for each feature distance, a first-order approximation can be written as

$$\sum_{j=1}^{3} \left(\frac{\partial \delta_i}{\partial T_j} \Delta T_j + \frac{\partial \delta_i}{\partial \phi_j} \Delta \phi_j\right) = \delta_i, \tag{15}$$

which can be rewritten as

$$\mathbf{J} \Delta \mathcal{P} = \delta, \tag{16}$$

where $\mathbf{J}$ is the Jacobian, $\mathcal{P}$ are the six parameters, and δ is the multidimensional error function. The method proceeds by computing new estimates for $\mathcal{R}$ and $\mathcal{T}$, and iterates until δ becomes acceptably small. The new estimates for $\mathcal{P}$ are obtained by inverting Eq. (16). Given the current estimate $\mathcal{R}_k$ and $\mathcal{T}_k$, and corrections $\Delta \mathcal{T}_k$ and $\Delta \phi_k$, the new estimates $\mathcal{R}_{k+1}$ and $\mathcal{T}_{k+1}$ can be computed as follows:

$$(1) \quad \begin{bmatrix} \mathcal{R}_{k+1} & \mathcal{T}_{k+1} \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} \mathcal{R}_k & \mathcal{T}_k \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \mathcal{R}_k & \Delta \mathcal{T}_k \\ 0 & 1 \end{bmatrix};$$

$$(2) \quad \mathcal{T}_{k+1} = \mathcal{T}_k + \Delta \mathcal{T}_k, \quad \phi_{k+1} = \phi_k + \Delta \phi_k;$$

$$(3) \quad \mathcal{T}_{k+1} = \mathcal{T}_k + \Delta \mathcal{T}_k, \quad \mathcal{R}_{k+1} = \mathcal{R}_k \Delta \mathcal{R}_k.$$

The third update formula is most robust, while the first two are not stable, as they are extremely sensitive to numerical errors. It is also important that the order of $\Delta \phi_1$, $\Delta \phi_2$, and $\Delta \phi_3$ should not be changed while generating $\Delta \mathcal{R}_k$. Alternately, some other representation for the rotation matrix can also be used.

5. Least-square robustness

The fiducial rotation is measured in radian while translation is measured in mm, giving rise to numerical instabilities during the inversion of the Jacobian in Eq. (16). The inversion process can further worsen if the Jacobian is close to singularity. Besides, the assignment of different weights to different features may be desired, since within the region of interest some features will be statistically more robust than others. These problems are solved by using parameter and task variable scaling,[39,40]

$$G(JH)(H^{-1}\Delta \mathcal{P}) = G\delta. \tag{17}$$

H scales the parameter space without changing the least-squares solution, since H can be chosen to make JH better conditioned than J. The H matrix can be constructed by using the parameter covariance matrix, for which we need an initial estimate of the solution and the standard deviation of each parameter. Column scaling[41] is another popular method when no *a priori* information is available. Unfortunately, this can have a problem if the parameters are poorly identifiable[42] in some poses. This is solved using the external scaling method, which does a column scaling on the most conservative estimates (the worst possible Jacobians). We ran 50 000 pose simulations to determine our H matrix using this method.

The G matrix is used for task variable (δ) scaling, which can be computed by the Gauss–Markov estimate by using a good estimate of the covariance matrix of the error function. In the absence of this matrix, another statistical approximation is to use a diagonal matrix, with each entry being the reciprocal of the standard deviation of that error element. This would make the L_2 norm a reasonable measure of the size of the least square error vector. We computed our G matrix using data from the same 50 000 pose simulations that we ran earlier.

6. Final optimization

The above framework was implemented using Newton's nonlinear optimization. We noticed that the optimization always converged in less than 25 iterations. The exit condition was defined as 50 iterations or a plateau in the error curve, whichever came later. We also incorporated some basic checks on the final answer to remove any mirror solutions. Other extraneous solutions can also act as a local minimum, in which case the initial estimate is perturbed a little to yield the correct convergence. A robust pose recovery algorithm was developed, by building on this optimization.
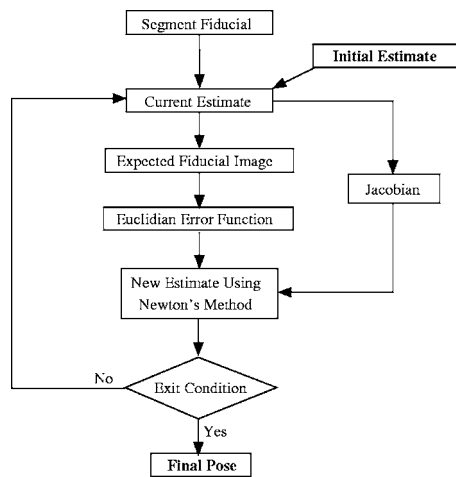
FIG. 5. The basic pose estimation algorithm.

FIG. 6. The robust pose estimation algorithm.

7. Robust pose recovery algorithm

To make the pose recovery more robust, we have designed a two-phase pose recovery algorithm. The basic pose recovery algorithm is outlined in Fig. 5. Like all optimizations, this one is prone to fall into a local minima, from which it may not be able to recuperate. This is resolved by running the algorithm multiple times from various initial estimates for the pose. This approach is able to reduce the probability of failure exponentially. In a loose sense, if $p(\sim 0.3)$ is the probability of failure in any given iteration, then the probability of success after N independent runs is $1 - p^N$, which is sufficiently close to guaranteed success. Moreover, to make pose recovery more accurate, we use the observation that point features, though inaccurate, converge more reliably (see Table I). Thus, the basic algorithm runs 25 times using only the point features, from which we obtain a very good guess for the final pose. This is used as the initial estimate and now the basic algorithm runs another 25 times using all features. Though the number of iterations can be cut to one-fifth without compromising on the convergence, the extremely fast run time allows us to add a safety bracket by increasing the number of iterations. Figure 6 outlines the algorithm. This robust algorithm ran for tens of thousands of poses (spanning the complete workspace) without failure.
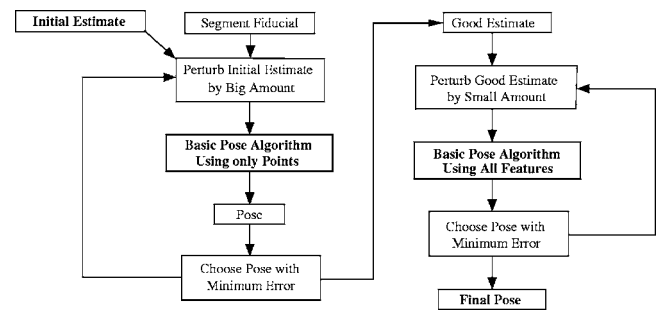
8. Segmentation

Additional constraints on their coordinates make a segmentation of parametric curves (straight lines, ellipses, etc.) more accurate than points. Since automatic segmentation of parametric curves is not the research contribution of this work, we do not produce specific numerical results that would sidetrack the main theme of the paper. Instead we refer the readers to the extensive literature already available in the computer vision and image processing community, where ellipse segmentation is a field in itself.[33,43–45] The parametric constraints can help segment the whole curve with high accuracy, even when most of the curve is occluded and only a small part is visible/segmented.[43] This makes pose estimation efficient. Moreover, the compact design of the fiducial decreases the likelihood of occlusion. We have developed and tested two methods of segmentation, both of which exhibit similar performances.

For many clinical applications, semiautomatic segmentation suffices. The operator clicks at a few points in the image that are close to the curves of interest. A minimum intensity centroid search in the neighborhood fine-tunes this location. Least square curve fitting is used to obtain the equation of the straight line and ellipse. Since the point features are known to be on the (extended) straight lines, their centroid positions are projected on the straight lines to get the final position. The semiautomatic MATLAB implementation takes about 10–15 s to process each image, which is acceptable for many applications.

Automatic segmentation of parametric curves is achieved by using the Hough transform,[33,44] which searches in the parametric space. An edge image is created to search in the parametric space. Straight lines are segmented using the Hough transform, after which each line is fine-tuned by doing a search for minimum intensity points in a small neighborhood, followed by running a RANSAC-based least square fitting[43] on these points. The three strongest near-parallel straight lines are chosen. The ellipse is composed of five parameters, which renders the parametric search space huge. A decoupled approach is used where the Hough transform is used to estimate various possible axes of symmetry in the image. Selected intersecting orthogonal axes are possible candidates for an ellipse. A 2D Hough transform is used to find the best ellipse for each intersection pair, which is then fine-tuned. The two strongest ellipses are chosen. Our imple-

TABLE I. For the basic pose algorithm, different feature combinations give varying errors and percentage failures.

Points, lines and ellipse combinations	Translation (nm)	Rotation (°)	Failures (%)
6P+3L+2E	1.82	0.89	31.40
6P+2E	1.96	1.16	28.00
6P+3L	2.09	1.24	28.80
3L+2E	2.23	0.98	90.60
6P	2.48	1.53	27.00

TABLE II. Pose recovery error (simulation) as a function of segmentation error.

Segmentation error (mm)	Translation (mm)			Rotation (°)		
	X	Y	Z	X	Y	Z
0.00	0.18	0.19	0.75	0.02	0.02	0.01
0.50	0.19	0.18	0.79	0.05	0.06	0.04
1.00	0.19	0.19	0.79	0.09	0.11	0.07
2.00	0.22	0.22	0.98	0.18	0.22	0.14
3.00	0.26	0.26	1.24	0.26	0.32	0.21

mentation is a normalized and robust variation of Ref. 45. The points are segmented as end points of the lines. Alternatively, a curve following technique[46,47] can also be used, though not implemented in our package.

III. RESULTS AND DISCUSSION

To measure sensitivity and robustness, the FTRAC fiducial was tested using simulated data and data that was obtained from real images of a precisely fabricated phantom. Finally, experiments measuring 3D fluoroscopic reconstruction of objects (*C*-arm tracked using the FTRAC fiducial) were conducted. The description and results for each experiment is followed immediately by the discussion.

A. Simulation studies

As a customary first step in the evaluation of fluoroscopic registration, simulations were conducted to evaluate the accuracy and robustness of the pose estimation software. Assuming an ideally calibrated fluoroscope, studies were conducted on synthetic images to analyse the effects of various governing parameters. MATLAB software was created to model x-ray imaging. Given *C*-arm parameters and fiducial pose, it generated a synthetic image and the exact location of each feature in the image. Appropriate noise was added to each imaging parameter, and its effect on pose recovery accuracy was observed. The error was modeled using a uniform and random probability density function, i.e., a 1 mm error means that a maximum error of *magnitude* 1 mm was added, with a uniform probability distribution. Based on our experience, the parameter values (for an average commercial *C* arm) were chosen to be 0.25 mm segmentation error, 1000±2 mm focal length, 0.5 mm translation from the image center for the origin, and 0.5 mm±1 μm as the pixel size. The region of use was near the isocenter of the *C* arm, at about a distance of two-thirds the focal length from the x-ray source. With the above-mentioned error in the parameters, individual parameters were changed to understand the sensitivity of pose recovery. The simulations were carried out for randomly selected poses in the workspace. The initial guess was also chosen at random, far away from the actual pose.

1. Segmentation

Segmenting the projected fiducial from the x-ray images is one of the most important sources of error. The last step of our segmentation algorithm is to pick out multiple minimum intensity points and fit a least square curve. Since in the synthetic data we know the precise position of each feature in the image, we added a segmentation error to the simulated point position. The erroneous points were then fit on a least square curve. This shifted the recovered position of the fiducial, which was compared to the known actual location.

Most of the translation error (Table II) was found to be along the imaging direction, which is intuitive since small changes in depth do not change the image significantly. The translation error parallel to the image plane is low, which is also intuitive since any small movement in that plane is magnified in the image. This can also be derived mathematically from Eq. (10), where the depth derivative has a quadratic term in the denominator, while the other derivatives are linear. Note that this could be an important source of error in depth sensitive surgeries, but in most cases is acceptable, especially when 3D information is projected into the image for guidance. The rotation error, on the other hand, is similar for all axes. The net translation and rotational error (based on 10 000 simulated poses) is displayed as a graph in Fig. 7(a). Most image processing algorithms achieve 1 pixel (0.4 mm) segmentation accuracy for ellipses, which indicates a 0.8 mm translation and 0.1° rotation error. This appears to be acceptable for most procedures and is similar to commercial image guidance systems that use optical/EM trackers. Moreover, the graph suggests that the pose recovery error does not grow exponentially with an increase in segmentation error. It can be concluded that the FTRAC fiducial is fairly robust to segmentation errors.

2. Focal length

The focal length, usually calibrated using an accurately machined phantom, could vary as much as a few mm when the *C* arm changes pose. Thus, an error in the focal length estimate is very likely, which could possibly lead to an inaccurate pose. Figure 7(b) (based on 20 000 simulated poses) shows the pose recovery error as a function of focal length calibration error. Orientation is always recovered with under 0.1° error, while translation error is nearly linear. Even a 5 mm error in focal length leads to under a 2 mm pose recovery error, with the largest component being in the imaging direction. We conclude that the orientation of the FTRAC fiducial is robust to focal length errors, while translation is a little sensitive.
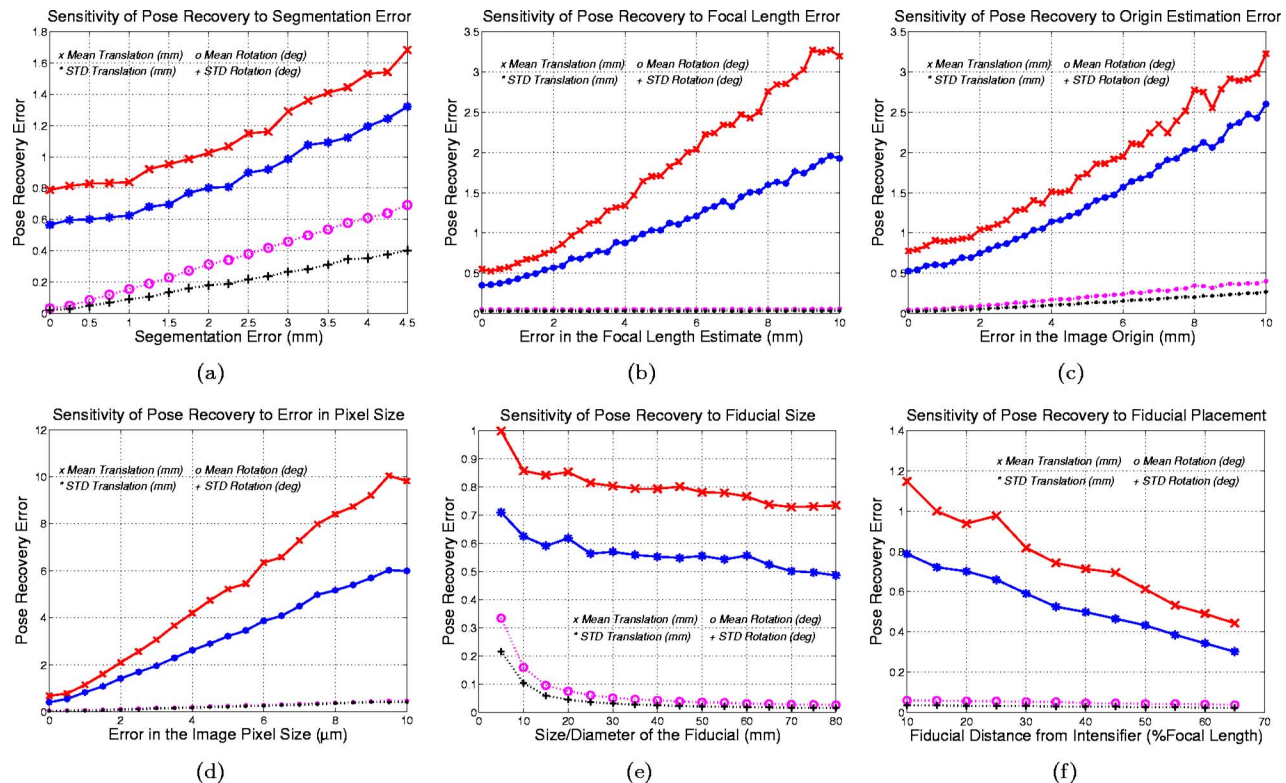
FIG. 7. Error in pose recovery as a function of (a) segmentation error; (b) focal length error; (c) origin error; (d) pixel size error; (e) fiducial size; (f) region of use.

3. Origin

The origin of the image, also calibrated using an accurately machined phantom, would change location with varying *C*-arm poses. Any error in the origin estimate could lead to a consistent shift in that particular direction. Figure 7(c) (based on 20 000 simulated poses) shows the rotational and translational sensitivity to the magnitude of error in origin estimation. Both translation and orientation errors are linear, though orientation is more robust. A 5 mm origin error leads to under a 2 mm translation error and 0.2° orientation error, with the largest component being parallel to the imaging direction. We conclude that the FTRAC fiducial is robust to origin estimation errors.

4. Pixel size

The pixel size along the *x* and *y* axes of the image does not change with fluoroscope movements and stays constant throughout the life of the fluoroscope. Hence a one-time calibration can provide very accurate estimates that can be used repeatedly, though any inaccuracies could have strong effects in fluoroscopic reconstruction. An error as small as 10 μm could lead to a 5 mm shift near the boundary of the image. Figure 7(d) (based on 20 000 simulated poses) shows the sensitivity to this parameter. Both translation and orientation errors grow linearly, with most of the translation error being parallel to the image and orientation error less than 0.5°. We conclude that pose recovery orientation is robust to errors in pixel size, while translation is sensitive but linear.

5. Fiducial size

In the design of the FTRAC fiducial, the size of the features can influence the accuracy of the algorithm to some extent. A large fiducial tends to be more stable than a smaller one, and thus the size should be chosen depending on the work space constraints of the clinical application. We studied the pose recovery sensitivity by scaling the size (diameter and length), but without altering the relative configuration. Figure 7(e) (based on 8000 simulated poses) shows the translation and rotation errors as a function of size. It can be concluded that the curve practically plateaus at 30 mm, with a bigger fiducial exhibiting no improvement in accuracy. Moreover, a fiducial of $1 \times 1 \times 2$ cm size appears to be practical for most applications, with only a marginal drop in accuracy.

6. Region of interest

The region of use of the fiducial also has a strong effect on the accuracies obtained. A fiducial closer to the source has a larger magnification, thus scaling any small errors and making the algorithm converge better. This is evident also from Eq. (10), where the denominator in the Jacobian is the depth of the fiducial. Figure 7(f) (based on 6000 simulated poses) shows the pose recovery error as a function of the distance from the image plane, expressed as a percentage of the focal length. It can be concluded that orientation is independent of fiducial placement, while the translation error drops linearly in depth.

TABLE III. Pose recovery results (phantom) using the FTRAC fiducial phantom.

Trial number	Number of images	Translation (mm)			Rotation (°)		
		X	Y	Z	X	Y	Z
1	4	0.03	0.04	0.40	0.13	0.07	0.07
2	7	0.13	0.04	0.68	0.21	0.27	0.27
3	7	0.04	0.02	0.51	0.15	0.15	0.12
4	7	0.05	0.02	0.52	0.17	0.17	0.11
5	13	0.08	0.10	0.60	0.40	0.40	0.15
Mean		0.07	0.04	0.55	0.21	0.21	0.15
STD		0.05	0.03	0.32	0.12	0.14	0.09

7. Run time and convergence

The robust pose recovery algorithm always converged without failure. An unoptimized implementation using MATLAB 6.5.13, on a 2.4 GHz Intel P4 with 512 MB of RAM with a Windows 2000 OS, takes around 3.5 s to run (excluding image transfer and processing time). This indicates that an optimized C++ implementation would perform near real time.

B. Pose recovery studies on phantom

A highly precise phantom was manufactured and real x-ray images were taken using a fluoroscope (Philips BV 3000). The system-supplied parameters were read from the DICOM header, otherwise the fluoroscope was not explicitly calibrated. Moreover, the images were not distortion corrected (distortion ≤ 2 mm). The fiducial was mounted on a highly accurate 0.002° resolution rotational turntable (30000 Heavy Load Worm Gear Drive from Parker Automation, Irwin, Pennsylvania). The fluoroscope remains stationary while the fiducial moves in a known path, providing ground truth.

A fiducial mount was designed such that it produced zero translation and a known rotation when the turntable was rotated. The design supported the two independent rotation axes typical to C arms. Thus, given any starting pose, the relative motion between the current pose and the starting pose is known precisely from the turntable reading. The relative motion is also calculated using our algorithm (from computed current pose and the computed starting pose). The difference between the computed relative motion and known relative motion is the error. This setup was taken to the OR and five series of tests were carried out with different fiducial trajectories. The results are shown in Table III.

We recorded a mean accuracy of 0.56 mm for translation (STD 0.33 mm) and 0.33° for rotation (STD 0.21°), which appears to be acceptable for most surgeries. It can be argued that there may be a constant drift in the pose algorithm computation, which would not be captured by this experiment. All five trials have consistent results with the simulation experiments, indicating that a constant error drift is highly unlikely, which is further affirmed by the 3D reconstruction studies below.

C. Reconstruction studies on phantom

C-arm tracking is not an end in itself, but is a necessary link toward 3D object reconstruction using multiple x-ray images. The second set of phantom experiments involved taking x-ray images of a stationary object from different views, and reconstructing it in space. The FTRAC fiducial was kept stationary, while the C arm was moved and multiple images (20° around the AP axis) were taken. The 3D coordinates of the nine fiducial points were reconstructed using two to five images (the FTRAC fiducial tracked the C arm). The location of these points was known with respect to the fiducial by precise fabrication. The 3D reconstruction accuracy is the difference between the computed coordinates and the actual coordinates of the fiducial points. The points were reconstructed with a mean accuracy of 0.53 mm (STD 0.16 mm), which appears to be acceptable in most surgeries. A significant portion of the pose error is in the imaging direction, and hence in order to improve accuracy, the image separation should be increased, being orthogonal in the ideal case. Complete results are available in Table IV. Note that this reconstruction accuracy will also be a function of image warping and segmentation error. The FTRAC fiducial has also been successfully used to reconstruct phantom brachytherapy implants,[48] the performance of which is similar to that reported here.

D. Comparison to conventional fiducials

In Sec. I, we provided the reasons why there is a need for a small, yet accurate radio-opaque fiducial. In this section, we offer quantitative results indicating progress over previous work. When analyzing the basic pose recovery algorithm, simulation experiments from Sec. II show that though using only point features is most stable, it is also the least accurate (Table I). As reviewed in Sec. I, all previously proposed fiducials can be classified as bead-based fiducials (except one[3] that we shall compare with in the end). Moreover, even among these fiducials, increasing the number of beads to more than 6 has little improvement in accuracy.[19] For any given fiducial, various pose estimation algorithms can be designed, most of which fall into two categories—linear and nonlinear. Linear algorithms evaluate least-squares solutions to the pose problem, for which they typically linearize the

TABLE IV. Object reconstruction results using the FTRAC fiducial phantom.

	Number of images used for each reconstruction	MEAN	STD	MIN	MAX
TRIAL 1	2	0.55	0.17	0.20	1.18
	3	0.50	0.14	0.19	0.92
	4	0.48	0.14	0.19	0.75
	5	0.47	0.14	0.23	0.67
	2	0.66	0.25	0.27	1.70
TRIAL 2	3	0.55	0.16	0.29	1.13
	4	0.53	0.15	0.30	0.80
	5	0.52	0.14	0.33	0.73

rotation matrices. Though quick and easy to implement, the problem with this approach is that it neglects the rigid rotation constraints, leading to lower accuracies. Extensive results are available in the literature,[20,23] where the accuracy of nonlinear methods is shown to be superior to linear methods. Thus, it should be sufficient to compare the performance of the FTRAC fiducial to a nonlinear pose estimation algorithm using six beads. To improve robustness of bead-based fiducials and to remove any biases that can arise due to a preferential configuration of the six beads, we compared the FTRAC fiducial to a nine-bead fiducial.

The predominant reason that the FTRAC fiducial offers better accuracies to other fiducials is that it relies on ellipses for evaluating the pose. Both for automatic and semiautomatic segmentation methods, achievable accuracies are dependent on multiple parameters like feature constraints, image resolution, x-ray wavelength, material used, choice of algorithm, etc. Thus, in the same image, different features provide different levels of accuracies. As reviewed in Sec. II, ellipses intrinsically offer better segmentation than beads. To quantitatively understand the improvement, we look at the residual errors from segmented datasets. The accuracy of each feature at the recovered pose is shown in Table V. It can be noticed that closed parametric curves (ellipses) segment with an overall accuracy of 0.13 mm, open parametric curves are a little more error prone with an accuracy of 0.19 mm, while points offer only 0.43 mm. The standard deviations

TABLE V. Residual distances for each feature from the five phantom trials. Parametric curves provide greater accuracies.

	Points	Lines	Ellipse
Residual error (mm)	0.44	0.14	0.15
	0.49	0.14	0.18
	0.43	0.28	0.11
	0.38	0.21	0.10
	0.42	0.19	0.11
MEAN	0.43	0.19	0.13
STD	0.04	0.05	0.03

values are stable at under one-eighth of a pixel (50 μm). Thus, it seems that ellipse segmentation tends to be about three times more accurate than points.

We conducted simulations for over 7000 poses and compared the pose recovery results of a nine-bead fiducial to the FTRAC fiducial, keeping this disparity in mind. The results for translational and rotational pose recovery accuracies are available in Figs. 8(a) and 8(b), respectively. It can be noted that even with high amounts of segmentation error, the FTRAC fiducial exhibits only a marginal loss in accuracy. On the other hand, the errors for bead-based fiducials climbs steadily. This pattern is observed for both translation and rotation, where the error quickly climbs to the 1 mm and 1° mark for the bead-based fiducials, while it remains consistently low for the FTRAC fiducial. This indicates that the FTRAC fiducial is not only more accurate, but also more robust when compared to conventional fiducials. We do not explicitly compare the FTRAC fiducial to the helical fiducial,[3] since the errors reported there are well over 1.5 mm in translation and 2° in rotation, while we obtain significantly superior results. Last, it should also be noted that the results obtained in Sec. III indicate a significant improvement to the previously published results (Sec. I).

IV. MOTIVATING APPLICATIONS

Though we have designed and tested the fiducial as a stand-alone piece of technology, we also intend to use it in a clinical application. The first use of this pose recovery technique will be in prostate brachytherapy, where *C* arms are ubiquitous, with over 60% of the practitioners using it in the operating room.[49] Unfortunately, it is used only for qualitative implant analysis and not for providing real-time intraoperative dosimetry. Thus, the ultimate goal will be to achieve registration between the seeds that are reconstructed from fluoroscopy and the soft tissue anatomy that is reconstructed from transrectal ultrasound (TRUS), which would allow us to make immediate provisions for any dose deviations from the intended plan. Several leading groups have published initial results favorably supporting *C*-arm fluoroscopy for intraoperative dosimetric analysis.[50–52] To achieve intraoperative dosimetry in prostate brachytherapy, we need
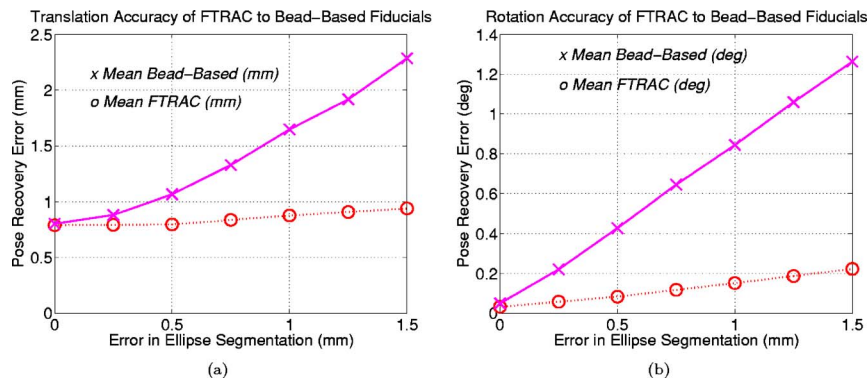
FIG. 8. A comparison of pose recovery accuracies for a nine-bead conventional fiducial compared to the FTRAC fiducial. Errors in (a) translation; (b) rotation. The FTRAC fiducial performs significantly superior.

to address the issues of (a) *C*-arm distortion correction and calibration; (b) *C*-arm pose tracking; (c) seed matching and reconstruction; and (d) registration of *C*-arm to TRUS images. The FTRAC fiducial is capable of addressing the issues of pose tracking and registration.

In a quick survey of registration, thin metal wire[53] inside a Foley catheter was used to visualize the prostatic urethra fluoroscopically in anterior–posterior and lateral projections. Recently, gold marker seeds[54] were implanted into the prostate and the relative positions of the needles and marker seeds were observed in fluoroscopy. Inserted implant needles[52] were used as fiducials for the registration of the fluoroscopy and ultrasound spaces. Unfortunately, implanted markers of any sort are susceptible to motion, and thus are not reliable. Alternately, radio-opaque spherical beads have been rigidly attached to the TRUS probe,[19] not only registering the TRUS to the x-ray images, but also addressing tracking. This process appears to require a permanent alteration to the probe, which for some practitioners might not be desirable.

Since in commercial brachytherapy systems the template is already registered to TRUS, the FTRAC fiducial will be attached to the implant template with a precision-machined mechanical connector in a known position. The spatial relationship between the FTRAC and TRUS is explained in Fig. 9. Let P_F be a point reconstructed in FTRAC space, P_U be the same point in ultrasound space, F_{FT} be the transformation between FTRAC and template, and F_{TU} be the transformation between template and ultrasound. Then the relationship between P_U and P_F can be expressed as $P_U = F_{TU} F_{FT} P_F$. Let us examine how accurate P_U is expected to be the accuracy of F_{FT} is defined by the accuracy of CNC machining and is in the neighborhood of 0.25 mm. The accuracy of P_F was found to be about 0.5 mm in our experiments. The accuracy of F_{TU} is specific to the FDA-approved brachytherapy system (CMS Interplant) that we will be using, and it is about 1.0 mm. Thus, the overall accuracy of the whole system is expected to be under 2 mm and can be further reduced by an improvement in the template-to-TRUS registration method. Experiments to evaluate the accuracy of the complete system are currently underway.

The fiducial is of general use, but has been optimized for prostate brachytherapy. The ellipses (and other features) are inclined at an angle of 30° to the vertical. Even when the

fiducial is kept close to the prostate and at varying heights, a 30° plane (to the vertical) divides the fiducial from the prostate (which contains the seeds). If the *C*-arm images are taken in any position inside the 30° cone around the AP axis, this design reduces the likelihood of overlap between the seed projections and the fiducial. As we have mentioned earlier, we will mount the FTRAC fiducial over the abdomen on the TRUS stepper, in a known calibrated pose relative to the template with the use of a precision-machined extender [Fig. 10(a)]. The distance between the FTRAC fiducial and the prostate will be about 15 cm, in which range the pose recovery and object reconstruction stabilities are adequate. The fallback plan,[55] as shown in Fig. 10(b), is to place the fiducial on a rectal sheath that supports the prostate, while the TRUS probe can move unhindered inside the sheath. The technical difficulty in this case was a robust coupling between the TRUS probe and the sheath, at acceptable loss in image quality. However, an increase in insertion diameter is of concern. A third alternative is to put the patient on a foam board incorporating one or several FTRAC fiducials in known relative offset with respect to the TRUS stepper.

Apart from brachytherapy, another promising application is fluoroguided robot-assisted surgery, as done in the literature.[3] If mounted on the robot, the FTRAC fiducial can register the robot to the *C*-arm space, achieving virtual fluoroscopy. In general, surgical tool tracking under x-ray fluo-
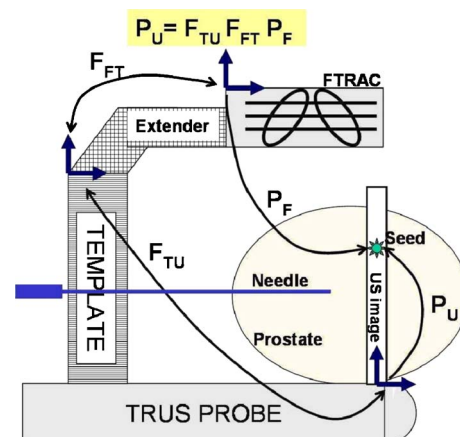

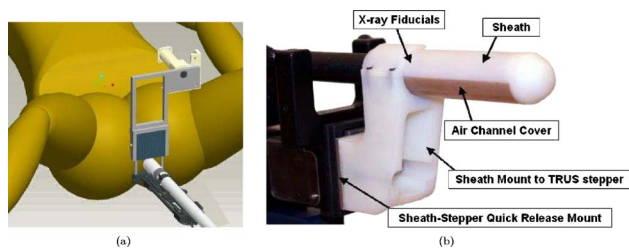FIG. 9. Frame transformations between the FTRAC and TRUS.

FIG. 10. The FTRAC fiducial mounted on (a) the template; (b) a hollow cylindrical sheath.

roscopy imaging can be achieved if the *C* arm is kept stationary, while the FTRAC fiducial is rigidly mounted on a surgical tool. Moreover, segmentation of the fiducial image and pose recovery have an inherent duality (i.e., knowing one directly relates to the other), which can further boost accuracy. The final pose can be used to improve segmentation, which, in turn, can enhance the pose. Iteratively doing this will not only improve accuracy, but also remove the need for an accurate segmentation. Alternately, an algorithm that does not optimize on the geometric error of the fiducial features, but optimizes directly on the intensity information, can be used. Both these techniques will eliminate the need for segmentation, while still achieving accurate pose recovery. The FTRAC fiducial can potentially also be used for both preoperative and intraoperative *C*-arm calibration. Additionally, the FTRAC fiducial can be instrumental in registering x-ray images to other imaging modalities. Thus, the FTRAC fiducial can facilitate a plurality of fluoroscopically guided procedures.

V. CONCLUSION AND OPEN QUESTIONS

A fiducial for robust and accurate *C*-arm tracking was designed, prototyped, computationally evaluated, and experimentally validated. The primary contribution of this work was the use of parametric curves (ellipses and straight lines), in addition to spherical beads that have been used conventionally. Point-based algorithms, though easy to implement, are sensitive to errors. On the other hand, parametric curves (a) segment accurately; (b) constrain pose recovery better; and (c) move pose recovery outside the framework of purely point-based methods. The proposed fiducial also addresses the issue of registration to other imaging modalities by precise placement.

Without distortion correction and extensive calibration, the FTRAC fiducial can track a C arm with a mean accuracy of 0.56 mm in translation (STD 0.33 mm) and 0.33° in rotation (STD 0.21°) and can reconstruct 3D coordinates to a mean accuracy of 0.53 mm (STD 0.16 mm). It not only promises to offer superior accuracies when compared to conventional fiducials, but also is significantly smaller in size ($3 \times 3 \times 5$ cm). The simulations suggest that the size of the FTRAC fiducial can be further reduced to $1 \times 1 \times 2$ cm, maintaining a similar level of accuracy.

One of the open questions in quantitative fluoroscopy is the importance of distortion correction and *C*-arm calibra-

tion. Generally speaking, distortion is widely prevalent and necessitates a correction procedure. This assures accurate tracking and reconstruction. Since properties vary from one machine to another, the distortion on some fluoroscopes might be small, while on others it might be large. Moreover, some applications might be more robust to distortion effects than others. Depending on the application and the amount of distortion, researchers might choose to (a) correct for a unique distortion at every pose, (b) use a constant correction for all poses, or (c) completely neglect distortion. This decision has to be made on a case by case basis. The results in this paper are on uncorrected images. The reason for robustness of our results to *C*-arm distortion can be attributed to the low amounts of distortion in the fluoroscope (average <2 mm) and to the small physical size of the fiducial, which covers only a small area in the image. Since distortion correction is not the focus of this paper, we do not analyze its effects on reconstruction. The results and analysis are available elsewhere.[48]

A similar issue arises for a calibration of the imaging parameters. Though it is believed that accurate calibration is a necessary prerequisite for quantitative fluoroscopy, our results indicate that it might not always be the case. Depending on the *C* arm used and the variation in the imaging parameters, inaccurate imaging parameters might still provide an accurate reconstruction of 3D information. These ideas can have desirable repercussions in a variety of clinical procedures, meriting further research. Though the FTRAC fiducial directly does not address these issues, it has provided the necessary intuition. We hope that these ideas, in collaboration with the FTRAC fiducial, will boost the clinical relevance and applicability of quantitative fluoroscopy on ordinary *C* arms.

ACKNOWLEDGMENTS

[a] Electronic mail: {jain,gabor}@cs.jhu.edu

[1] R. Hofstetter, M. Slomczykowski, M. Sati, and L. Nolte, "Fluoroscopy as an imaging means for computer-assisted surgical navigation," Comput. Aided Surg. **4**, 65–76 (1999).

[2] L. Joskowicz, C. Milgrom, A. Simkin, L. Tockus, and Z. Yaniv, "Fracas: a system for computer-aided image-guided long bone fracture surgery," Comput. Aided Surg. **3**, 271–288 (1998).

[3] J. Yao, R. H. Taylor, R. P. Goldberg, R. Kumar, A. Bzostek, V. R. Van, P. Kazanzides, and A. Gueziec, "A c-arm fluoroscopy-guided progressive cut refinement strategy using a surgical robot," Comput. Aided Surg. **5**, 373–390 (2000).

[4] J. C. Eyke, J. E. Ricciardi, W. Roesch, and T. S. Whitecloud, "Computer-assisted virtual fluoroscopy," University of Pennsylvania Orthopaedic Journal, 2002, pp. 53–59.

[5] K. Wiesent, K. Barth, N. Navab, P. Durlak, T. Brunner, O. Schuetz, and W. Seissler, "Enhanced 3-d-reconstruction algorithm for c-arm systems suitable for interventional procedures," IEEE Trans. Med. Imaging **19**, 412–422 (2000).

[6] S. A. M. Baert, W. J. Niessen, and M. A. Viergever, "Guide wire tracking during endovascular interventions," IEEE Trans. Med. Imaging **22**, 965–972 (2003).

[7] T. Yamazaki, T. Watanabe, Y. Nakajima, K. Sugamoto, T. Tomita, H. Yoshikawa, and S. Tamura, "Improvement of depth position in 2-d/3-d registration of knee implants using single-plane fluoroscopy," IEEE Trans. Med. Imaging **23**, 602–612 (2004).

[8] T. Rohlfing, D. Russako, M. Murphy, and C. Maurer, "An intensity-based registration algorithm for probabilistic images and its application for 2-d to 3-d image registration," in *SPIE Medical Imaging: Image Processing*, 2002, Vol. 4684, pp. 581–591.

[9] T. Tang, "Calibration and point-based registration of fluoroscope images," Master's thesis, Queen's University, January 1999.

[10] R. Siddon and N. Barth, "Stereotaxic localization of intracranial targets," Int. J. Radiat. Oncol., Biol., Phys. **13**, 1241–1246, August 1987.

[11] R. Koppe, E. Klotz, J. O. d. Beek, and H. Aerts, "3d vessel reconstruction based on rotational angiography," in *CAR*, 1995, pp. 101–107.

[12] N. Navab, A. Bani-Hashemi, M. Mitschke, D. W. Holdsworth, R. Fahrig, A. J. Fox, and R. Graumann, "Dynamic geometrical calibration for 3-d cerebral angiography," in *SPIE Medical Imaging*, February 1996, Vol. 2708, pp. 361–370.

[13] C. Brack, M. Roth, A. Czopf, J. Moctezuma, F. Gtte, and A. Schweikard, "Towards accurate x-ray camera calibration in computer assisted robotic surgery," in *Computer-Assisted Radiology and Surgery (CARS)*, Berlin, 1996, pp. 721–728.

[14] M. Mitschke and N. Navab, "Recovering projection geometry: How, a cheap camera can outperform an expensive stereo system," in *CVPR*, 2001, Vol. 1, pp. 193–200. Available online at citeseer.ist.psu.edu/456732.html

[15] VectorVision® Navigation System, "Brainlab, Inc., Heimstetten, Germany."

[16] StealtStation®, "Medtronic Surgical Navigation Technologies, Louisville, CO."

[17] OEC 9800 FluoroTrak™, "GE Healthcare, Waukesha, WI."

[18] R. Phillips, A. Mohsen, W. Viant, S. Malek, Q. Li, N. Shah, M. Bielby, and K. Sherman, "A phantom based approach to fluoroscopic navigation for orthopaedic surgery," in *MICCAI (2)*, 2004, pp. 621–628.

[19] M. Zhang, M. Zaider, M. Worman, and G. Cohen, "On the question of 3d seed reconstruction in prostate brachytherapy: the determination of x-ray source and film locations," Phys. Med. Biol. **49**, N335–345 (2004).

[20] Z. Yaniv and L. Joskowicz, "Precise robot-assisted guide positioning for distal locking of intramedullary nails," IEEE Trans. Med. Imaging **24**, 624–635 (2005).

[21] A. K. Jain, T. Mustufa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (FTRAC) fiducial," in *SPIE Medical Imaging; Visualization, Image-Guided Procedures, and Display*, February 2005, Vol. 5744, pp. 798–809.

[22] T. S. Y. Tang, N. J. MacIntyre, H. S. Gill, R. A. Fellows, N. A. Hill, D. R. Wilson, and R. E. Ellis, "Accurate assessment of patellar tracking using fiducial and intensity-based fluoroscopic techniques," Med. Image Anal **8**, 343–351 (2004).

[23] D. DeMenthon and L. Davis, "Model-based object pose in 25 lines of code," Int. J. Comput. Vis. **15**, 123–141 (1995).

[24] J. Fitzpatrick and J. West, "The distribution of target registration error in rigid-body point-based registration," IEEE Trans. Med. Imaging **20**, 917–927 (2001).

[25] J. Heikkila and O. Silven, "A four-step camera calibration procedure with implicit image correction," in *CVPR97*, 1997, pp. 1106–1112.

[26] A. K. Jain, T. Mustufa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (FTRAC) fiducial," in *AAPM Annual Meeting*, Pittsburgh, July 2004.

[27] O. D. Faugeras and B. Mourrain, "On the geometry and algebra of the point and line correspondences between *n* images," in *ICCV*, 1995, pp. 951–956. [Online]. Available: citeseer.ist.psu.edu/faugeras95geometry.html

[28] A. Ansar and K. Daniilidis, "Linear pose estimation from points or lines," in *ECCV02*, 2002, Sec. IV, pp. 282ff.

[29] Q. Ji, M. Costa, R. Haralick, and L. Shapiro, "An integrated linear technique for pose estimation from different features," *International Journal of Pattern Recognition and Artificial Intelligence*, June 1999. Available online at citeseer.ist.psu.edu/ji99integrated.html

[30] C. Taylor and D. Kriegman, "Structure and motion from line segments in multiple images," *T-PAMI*, 1995, Vol. 17, pp. 1021–1032.

[31] R. Hu and Q. Ji, "Camera self-calibration from ellipse correspondences," in *ICRA*, 2001, pp. 2191–2196.

[32] L. Wang, C. Grimm, and R. Pless, "A 3d pattern for post estimation for object capture," Available online at citeseer.ist.psu.edu/606684.html

[33] E. Trucco and A. Verri, *Introductory Techniques for 3-D Computer Vision* (Prentice–Hall, Englewood Cliffs, NJ, 1998), pp. 283–294.

[34] P. Yao, G. Evans, and A. Calway, "Face tracking and pose estimation using affine motion parameters," in *Proceedings of the 12th Scandinavian Conference on Image Analysis*, I. Austvoll, Ed. Norwegian Society for Image Processing and Pattern Recognition, 2001, pp. 531–536. [Online]. Available: citeseer.ist.psu.edu/yao01face.html

[35] M. Dhome, J. La Preste, G. Rives, and M. Richetin, "Spatial localization of modelled objects of revolution in monocular perspective vision," in *ECCV90*, 1990, pp. 475–485.

[36] S. Ma, "Conics-based stereo, motion estimation and pose determination," *IJCV*, February 1993, Vol. 10, No. 1, pp. 7–25.

[37] M. Costa and L. Shapiro, "3d object recognition and pose with relational indexing," *CVIU*, September 2000, Vol. 79, No. 3, pp. 364–407.

[38] D. Eberly, "Distance from a point to an ellipse in 2D," Oct 2002, www.magic-software.com, Magic Software, Inc. Available Online at www.magic-software.com

[39] J. Hollerbach and C. Wampler, "The calibration index and taxonomy for robot kinematic calibration methods," *IJRR*, 1996, Vol. 15(6), pp. 573–591, 1996.

[40] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (Prentice–Hall, Englewood Cliffs, NJ, 1974).

[41] R. Prager, R. Rohling, A. Gee, and L. Berman, "Automatic calibration for 3-d free-hand ultrasound," Cambridge University Department of Engineering, Tech. Rep. CUED/F-INFENG/TR 303, September 1997.

[42] K. Schroer, *Theory of Kinematic Modeling and Numerical Procedures For Robot Calibration* (Chapman & Hall, London, 1993), pp. 157–196.

[43] R. Halir and J. Flusser, "Numerically stable direct least squares fitting of ellipses," available online at citeseer.ist.psu.edu/350661.html

[44] R. McLaughlin, "Randomized Hough transform: Improved ellipse detection with comparison," Pattern Recogn. Lett. **19**, 299–305 (1998). Available online at citeseer.ist.psu.edu/mclaughlin98randomized.html

[45] Y. Lei and K. C. Wong, "Ellipse detection based on symmetry," Pattern Recogn. Lett. **21**, 41–47 (1999).

[46] C.-L. T. Tsai, C. V. Stewart, H. L. Tanenbaum, and B. Roysam, "Model-based method for improving the accuracy and repeatability of estimating vascular bifurcations and crossovers from retinal fundus images," IEEE Trans. Inf. Technol. Biomed. **8**, 122–130 (2004).

[47] A. Can, H. Shen, J. N. Turner, H. L. Tanenbaum, and B. Roysam, "Automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms," IEEE Trans. Inf. Technol. Biomed. **3**, 125–138 (1999).

[48] A. K. Jain, Y. Zhou, T. Mustufa, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)," in *SPIE Medical Imaging; Visualization, Image-Guided Procedures, and Display*, February 2005, Vol. 5744, pp. 810–821.

[49] B. Prestidge, J. Prete, T. Buchholz, J. Friedland, R. Stock, P. Grimm, and W. Bice, "A survey of current clinical practice of permanent prostate brachytherapy in the United States," Int. J. Radiat. Oncol., Biol., Phys. **15**, 461–465 (1998).

[50] D. Todor, G. Cohen, H. Amols, and M. Zaider, "Operator-free, film-based 3d seed reconstruction in brachytherapy," Phys. Med. Biol. **47**, 2031–2048 (2002).

[51] D. Todor, M. Zaider, G. Cohen, M. Worman, and M. Zelefsky, "Intraoperative dynamic dosimetry for prostate implants," Phys. Med. Biol. **48**, 1153–1171 (2003).

[52] L. Gong, P. Cho, B. Han, K. Wallner, S. Sutlief, S. Pathak, D. Haynor, and Y. Kim, "Ultrasonography and fluoroscopic fusion for prostate brachytherapy dosimetry," Int. J. Radiat. Oncol., Biol., Phys. **1**, 1322–1330 (2002).

[53] K. Wallner, J. Roy, M. Zelefsky, Z. Fuks, and L. Harrison, "Fluoroscopic visualization of the prostatic urethra to guide transperineal prostate implantation," Int. J. Radiat. Oncol., Biol., Phys. **1**, 863–867 (1994).

[54] M. Baird, R. Holt, and T. Selby, "Improvement of transperineal implant dosimetry by intraoperative cystoscopic confirmation of prostate anatomy," J. Urol. (Paris) **164**, 406–410 (2000).

[55] J. Kemper, A. Burkholder, A. Jain, T. Mustufa, K. Wyrobek, C. Burdette, D. Song, A. Okamura, and G. Fichtinger, "Transrectal fiducial carrier for radiographic image registration in prostate brachytherapy," in *AAPM Annual Meeting*, Seattle, July 2005.